z/OS

Parallel Sysplex Overview: Introducing Data Sharing and Parallelism in a Sysplex

SA22-7661-00

z/OS

Parallel Sysplex Overview: Introducing Data Sharing and Parallelism in a Sysplex

SA22-7661-00

Note

Before using this information and the product it supports, be sure to read the general information under “Appendix. Notices” on page 41.

First Edition, March 2001

This edition applies to Version 1 Release 1 of z/OS (5694-A01), and to subsequent releases and modifications until otherwise indicated in new editions.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

IBM welcomes your comments. A form for readers’ comments may be provided at the back of this publication, or you may address your comments to the following address:

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400
United States of America

FAX (United States & Canada): 1+845+432-9405
FAX (Other Countries):
Your International Access Code +1+845+432-9405

IBMLink (United States customers only): IBMUSM10(MHVRCFS)
Internet e-mail: mhvrcfs@us.ibm.com
World Wide Web: http://www.ibm.com/servers/eserver/zseries/zos/webqs.html

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:
- Title and order number of this book
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Figures

About This Book

This publication is about data sharing and parallelism in a sysplex and how it can meet the needs of your computing environment.

It introduces high-level concepts and information useful for anyone who will plan for, exploit, implement, or operate a sysplex that shares data and processes work in parallel.

This publication describes:
- What a sysplex is and how it compares to existing non-sysplex computing environments
- New features that enhance sysplex function
- Work that a sysplex supports
- Products that make up a sysplex
- Various sysplex configurations and their benefits

Sysplex and Product Availability (“Roll-Out”) The sysplex is a large system computing environment that is evolving. Since the introduction of the sysplex, the coupling facility technology was developed to enhance sysplex capabilities. With a coupling facility in a sysplex, the participating MVS systems can do high performance data sharing. A sysplex with a coupling facility is called a ***Parallel Sysplex***.

Note that this publication does not differentiate between a sysplex without a coupling facility and a sysplex with a coupling facility. When you see the term *sysplex*, understand it to mean a sysplex with a coupling facility (a Parallel Sysplex).

There might be changes to the implementation or availability of new products or functions. For information about availability of hardware and software sysplex support, consult your IBM representative.

Who Should Use This Publication

This publication is intended for anyone who wants an introduction to data sharing and parallelism in a sysplex. It discusses sysplex evolution, purpose, benefits, products, and services (hardware and software). The reader can be an executive, technical consultant, information systems manager, technical planner, or system programmer.

Notes on Terminology

Please note the following regarding terminology in this publication:
- When you see the term ***Parallel Sysplex***, understand it to mean a sysplex with a coupling facility, which is a ***Parallel Sysplex***.
- When you see the term ***MVS*** in this book, understand it to mean the element of either OS/390 or z/OS.
- The first time a product is referenced in the text, its entire name is spelled out and its acronym specified. From then on, only the acronym is used. For quick reference, the glossary contains product acronyms used in this publication, and what they stand for.

- The term ***system*** is one that is used heavily in many publications to mean many different things. Often you must attempt to discern its exact meaning from the context in which it is used. This publication uses the term ***system*** to mean an ***MVS system***. Here are some important definitions:

  MVS system: An MVS image together with its associated hardware, which collectively are often referred to simply as a system, or MVS system.

  MVS image: A single occurrence of the operating system that has the ability to process work.
- In this publication, the licensed program DB2 for OS/390 is referred to as ***DB2***.
- In this publication, the licensed program CICS Transaction Server for OS/390 is called **CICS Transaction Server** and refers to the product that contains Customer Information Control System (CICS), CICSPlex Systems Management (CICSPLex SM), and others.
- The licensed program System Automation for OS/390 includes all of the function previously provided by the program products Automated Operations Control/MVS (AOC/MVS), Target System Control Facility (TSCF), and ESCON Manager.
- In the figures, the term S/390 microprocessor complex or microprocessor cluster means the **S/390 9672 Parallel Transaction Server, 9672 Parallel Enterprise Server models, or IBM zSeries 900.**

Notes on Product Availability

The term “follow-on phase” will be used whenever discussing function planned to be available in future phases of the roll-out of sysplex support across products. If “follow-on phase” is not specified, assume that the function is available or announced. Note that the information about follow-on phases represents IBM’s **intent**, and is subject to change or withdrawal.

Where to Find More Information

This publication is part of a library of Parallel Sysplex planning books. IBM strongly recommends that you read each of the books in this library before you begin planning your Parallel Sysplex. Here is the complete list of publications in the Parallel Sysplex planning library, along with their order numbers:

Table 1. Parallel Sysplex Library Books

Publication Title	Order Number	Description
z/OS Parallel Sysplex Overview	SA22-7661	Use this publication as an overview to the sysplex and coupling facility data sharing. It describes highlights and the value of the sysplex to your business.
z/OS Parallel Sysplex Application Migration	SC22-7662	Use this publication to understand planning considerations for the parallel processing of applications that run on Customer Information Control System (CICS)/ESA and Information Management System (IMS)/ESA software in the Parallel Sysplex.

Using LookAt to look up message explanations

LookAt is an online facility that allows you to look up explanations for z/OS messages and system abends.

Using LookAt to find information is faster than a conventional search because LookAt goes directly to the explanation.

LookAt can be accessed from the Internet or from a TSO command line.

You can use LookAt on the Internet at:

`http://www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/lookat.html`

To use LookAt as a TSO command, LookAt must be installed on your host system. You can obtain the LookAt code for TSO from the LookAt Web site by clicking on **News and Help** or from the *z/OS Collection*, SK3T-4269.

To find a message explanation from a TSO command line, simply enter: **lookat** *message-id* as in the following example:

```
lookat iec192i
```

This results in direct access to the message explanation for message IEC192I.

To find a message explanation from the LookAt Web site, simply enter the message ID. You can select the release if needed.

Note: Some messages have information in more than one book. For example, IEC192I has routing and descriptor codes listed in *z/OS MVS Routing and Descriptor Codes*. For such messages, LookAt prompts you to choose which book to open.

Accessing licensed books on the Web

z/OS licensed documentation in PDF format is available on the Internet at the IBM Resource Link Web site at:

`http://www.ibm.com/servers/resourcelink`

Licensed books are available only to customers with a z/OS license. Access to these books requires an IBM Resource Link Web userid and password, and a key code. With your z/OS order you received a memo that includes this key code.

To obtain your IBM Resource Link Web userid and password log on to:

`http://www.ibm.com/servers/resourcelink`

To register for access to the z/OS licensed books:

1. Log on to Resource Link using your Resource Link userid and password.
2. Click on **User Profiles** located on the left-hand navigation bar.
3. Click on **Access Profile.**
4. Click on **Request Access to Licensed books.**
5. Supply your key code where requested and click on the **Submit** button.

If you supplied the correct key code you will receive confirmation that your request is being processed. After your request is processed you will receive an e-mail confirmation.

Note: You cannot access the z/OS licensed books unless you have registered for access to them and received an e-mail confirmation informing you that your request has been processed.

To access the licensed books:

1. Log on to Resource Link using your Resource Link userid and password.
2. Click on **Library**.
3. Click on **zSeries**.

4. Click on **Software**.
5. Click on **z/OS**.
6. Access the licensed book by selecting the appropriate element.

Chapter 1. Introduction to Sysplex

The word ***sysplex*** might not be new to you. The MVS/ESA sysplex has been available since 1990 when it was announced as a platform for an evolving large system computing environment.

What's new is what a sysplex can now do for you. It's more than a platform. It has become the large system computing environment that offers you improved price/performance through cost effective processor technologyand enhanced software. This technology builds on existing data processing skills and will run existing applications—an additional cost saver. A sysplex can also increase system availability, and at the same time, it increases your potential for doing more work.

In case you're not sure what a sysplex is, this chapter starts with a definition of a sysplex, and then describes the new capabilities of a sysplex and why you might be interested in it.

What is a Sysplex?

A sysplex is a collection of MVS systems that cooperate, using certain hardware and software products, to process work. A conventional large computer system also uses hardware and software products that cooperate to process work. A major difference between a sysplex and a conventional large computer system is the improved growth potential and level of availability in a sysplex. The sysplex increases the number of processing units and MVS operating systems that can cooperate, which in turn increases the amount of work that can be processed. To facilitate this cooperation, new products were created and old products were enhanced.

The following diagram shows the visible parts of a sysplex, namely the hardware. The software, which helps to enable the interaction and cooperation among the systems, is described later.

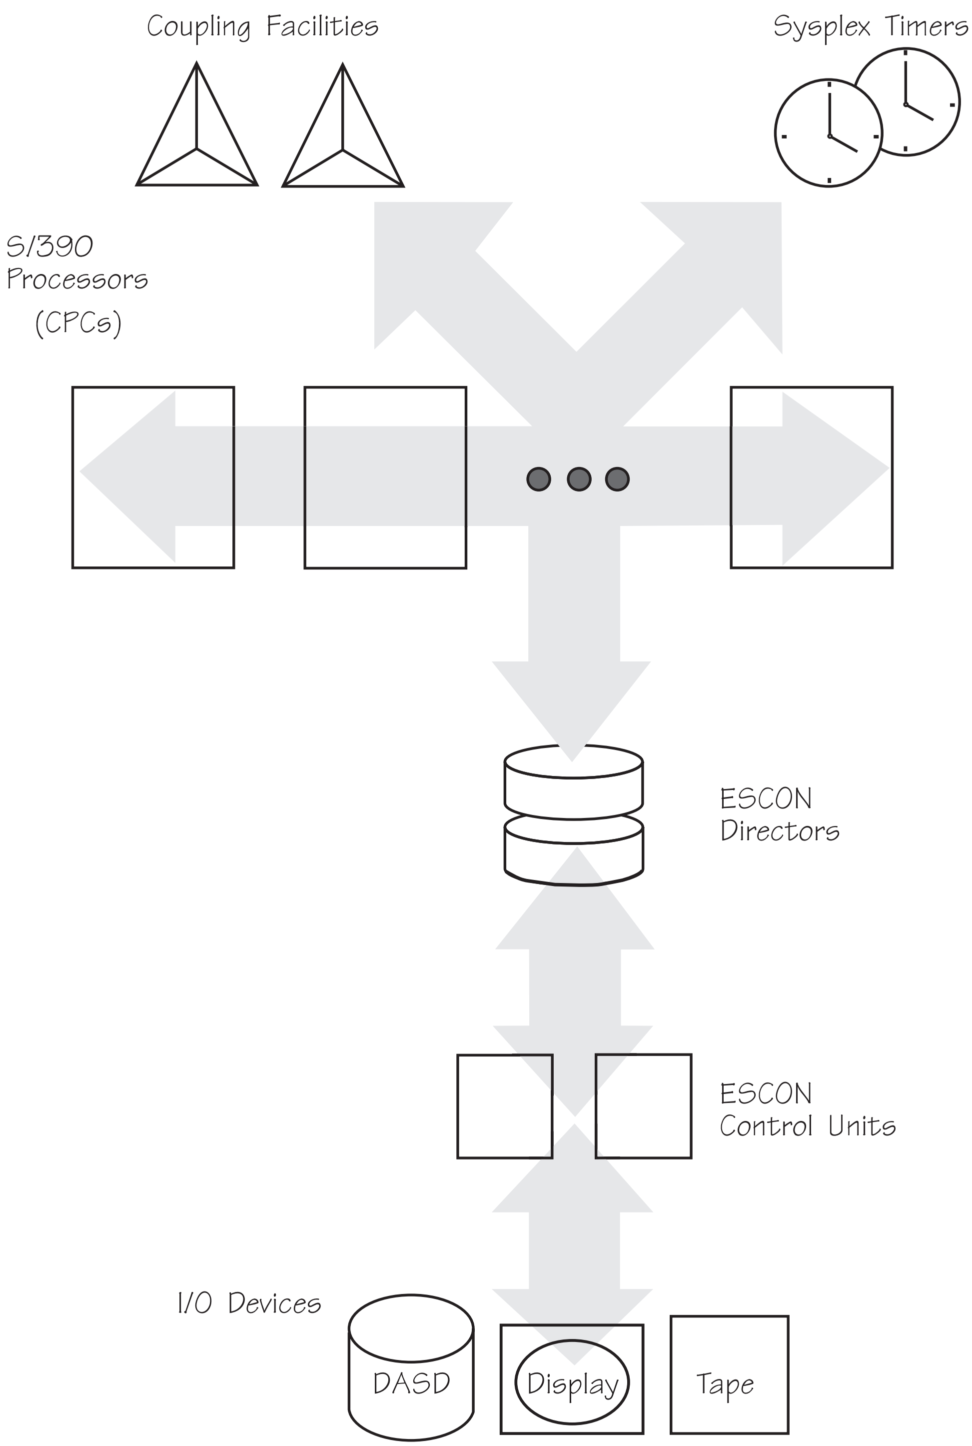

Figure 1. Sysplex Structure

Sysplex Hardware

The following types of hardware participate in a sysplex.

Coupling Facility: The coupling facility enables high performance multisystem data sharing and includes the S/390 9674 Coupling Facility. You can also define logical partitions in the S/390 9672 Parallel Sysplex Transaction Server or 9672 Parallel Sysplex Enterprise Server as coupling facilities.

Coupling Facility Channels: Coupling facility provides high speed connectivity between the coupling facility and the central processor complexes (CPCs) that use it.

Sysplex Timers: A Sysplex Timer synchronizes the time-of-day (TOD) clocks in multiple CPCs in a sysplex.

Processors: Selected models of S/390 or z900 processors can take advantage of a sysplex. These include large water-cooled processors, air-cooled processors, and the processors that take advantage of (CMOS) technology including the 9672 Parallel Transaction Server, 9672 Parallel Enterprise Server R1/R2/R3 Models.

ESCON Channels and Directors: Enterprise Systems Connection (ESCON) channels enhance data access and communication in the sysplex. The ESCON Directors add dynamic switching capability for ESCON channels.

ESCON Control Units and I/O Devices: ESCON control units and I/O devices in a sysplex provide the increased connectivity necessary among a greater number of systems.

Sysplex Software

The following software works with the hardware to enable interaction and cooperation among systems in a sysplex.

System Software: Base system software that is enhanced to support a sysplex includes the MVS operating system, JES2 and JES3, and DFSMS).

Networking Software: Virtual Telecommunications Access Method (VTAM) supports the attachment of a sysplex to a network.

Data Management Software: The data managers that support data sharing in a sysplex are Information Management System Database Manager (IMS DB), DATABASE 2 (DB2), and Virtual Storage Access Method (VSAM).

Transaction Management Software: The transaction managers that support a sysplex are Customer Information Control System (CICS/ESA and CICS Transaction Server) and Information Management System Transaction Manager (IMS TM).

Systems Management Software: A number of software products are enhanced to run in a sysplex and exploit its capabilities. The products manage accounting, workload, operations, performance, security, and configuration, and they make a sysplex easier to manage by providing a single point of control.

More information about the specific hardware and software products is available in “Chapter 4. Products in a Sysplex” on page 31.

Why Sysplex?

Now that you've been introduced to the pieces that make up a sysplex, you might be wondering what a sysplex could do for you. If your data center is responsible for even one of the following types of work, you could benefit from a sysplex.

- Large business problems—ones that involve hundreds of end users, or deal with a very large volume of work that can be counted in the millions of transactions per day.
- Work that consists of small work units, such as online transactions, or large work units that can be subdivided into smaller work units, such as queries.
- Concurrent applications on different systems that need to directly access and update a single database without jeopardizing data integrity and security.

A sysplex shares the processing of work across MVS/ESA systems, and as a result offers benefits, such as:

- Reduced cost through:
  - Cost effective processor technology
  - Continued use of large system data processing skills without re-education
  - Protection of application investments
  - The ability to manage a large number of systems more easily than other comparably performing multisystem environments
- Platform for continuous availability so that applications can be available 24 hours a day, 7 days a week, 365 days a year (or close to it)
- Ability to do more work
  - Greater capacity
  - Improved ability to manage response time
  - Platform for further capacity and response time advances
- Greater flexibility
  - Ability to mix levels of hardware and software
  - Ability to dynamically add systems
  - An easy path for incremental growth
  - Varied platforms for applications, including parallel, open, and client/server

Depending on your data center's goals and needs, some of these benefits might be more attractive to you than others. Your IBM marketing representative will be able to discuss, in more detail, how a benefit can specifically apply to your situation.

An Evolutionary Approach

A sysplex is the most recent development in the evolution of IBM large systems. Large system configurations have evolved from a single system uniprocessor through tightly coupled multiprocessors, to a loosely coupled configuration, to the sysplex.

Single System Uniprocessor

A system is made up of hardware products including a central processor (CP), and software products, with the primary one being an operating system such as MVS. Other types of software—system application programs, end-user application programs, tools—run on the system. The CP is the functional hardware unit that interprets and processes program instructions. The CP and other system hardware, such as channels and storage, make up a central processor complex (CPC).

The System/390 architecture defines that a single CP process one and only one instruction from a program at a time. The MVS operating system manages the instructions to be processed and the resources required to process them. When a single copy of the MVS operating system (MVS image) manages the processing of a CPC that has a single CP, the system configuration is called a uniprocessor.

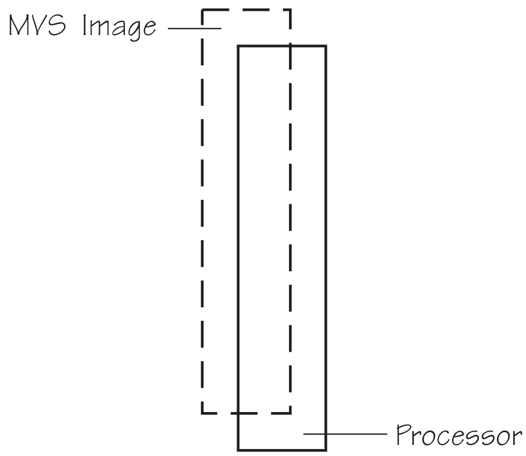

Figure 2. A Uniprocessor with One MVS Image

Summary

Capacity
Equal to the size of the largest single CP

Availability
Contains single points of failure; disruptive change

Systems Management
Easy to manage work

Tightly Coupled Multiprocessors

When you add more CPs to the central processor complex (CPC), you add the capability of processing program instructions simultaneously. When all the CPs share central storage and a single MVS image manages the processing, work is assigned to a CP that is available to do the work. If a CP fails, work can be routed to another CP. This hardware and software organization is called a tightly coupled multiprocessor.

A tightly coupled multiprocessor has more than one CP and a single MVS image sharing central storage. The CPs are managed by the single MVS image, which assigns work to them.

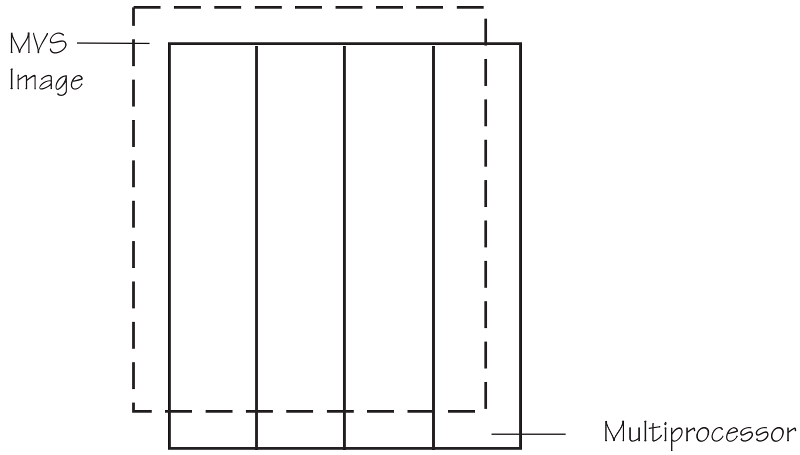

Figure 3. A Tightly Coupled Multiprocessor

Summary

Capacity
Increased over that of a uniprocessor but limited by the maximum number of CPs in the CPC

Availability
Increased over that of a uniprocessor but limited by some characteristics, including one system image, that may represent single points of failure and disruptive change

Systems Management
Easy to manage work

Loosely Coupled Configuration

A tightly coupled multiprocessor provides CP backup in case of failure. But what about software backup? Can you have more than one MVS image in a configuration?

Systems outside a sysplex can coordinate more than one MVS image with the MVS job entry subsystem components (JES2 and JES3) and global resource serialization. These components establish a means of sharing a work input queue across a number of systems to allow shared data sets, printers, and consoles. This type of system configuration is called loosely coupled.

A loosely coupled configuration has more than one CPC, possibly tightly coupled multiprocessors, sharing DASD but not central storage. The CPCs can connect by channel-to-channel communications and are managed by more than one MVS image. Work is distributed from a shared job queue to each MVS.

Although a loosely coupled configuration increases system capacity, it is not as easy to manage as either a uniprocessor or a tightly coupled multiprocessor. Each system must be managed separately, often by a human operator, who monitors product-specific messages on a set of consoles for each system.

Products and applications that need to communicate and are running on separate systems have to create their own communication mechanism. These varied communication mechanisms add to the difficulty of managing a loosely coupled configuration.

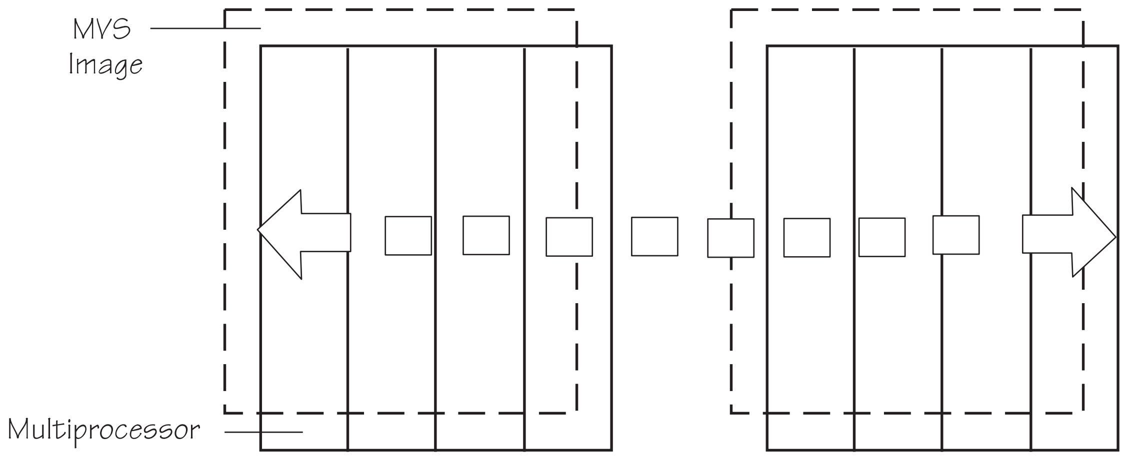

Figure 4. A Loosely Coupled Configuration

Summary

Capacity
Increased system capacity over tightly coupled multiprocessors and a uniprocessor

Availability
Increased system availability over tightly coupled multiprocessors and a uniprocessor

Systems Management
Requires additional systems management—separate MVS images communicate to share data sets, printers, and consoles

Base Sysplex

To help solve the difficulties of managing many MVS systems, IBM introduced the MVS ***sys***tems com***plex*** or sysplex in September of 1990. The base sysplex lays the groundwork for simplified multisystem management through the cross-system coupling facility (XCF) component of MVS/ESA. XCF services allow authorized applications on one system to communicate with applications on the same system or on other systems. In a base sysplex, CPCs connect by channel-to-channel communications and a shared dataset to support the communication. When more than one CPC is involved, a Sysplex Timer synchronizes the time on all systems.

The base sysplex is similar to a loosely coupled configuration in that more than one CPC (possibly a tightly coupled multiprocessor) shares DASD and is managed by more than one MVS image. A sysplex is different from a loosely coupled configuration because through XCF, there is a standard communication mechanism for MVS system applications.

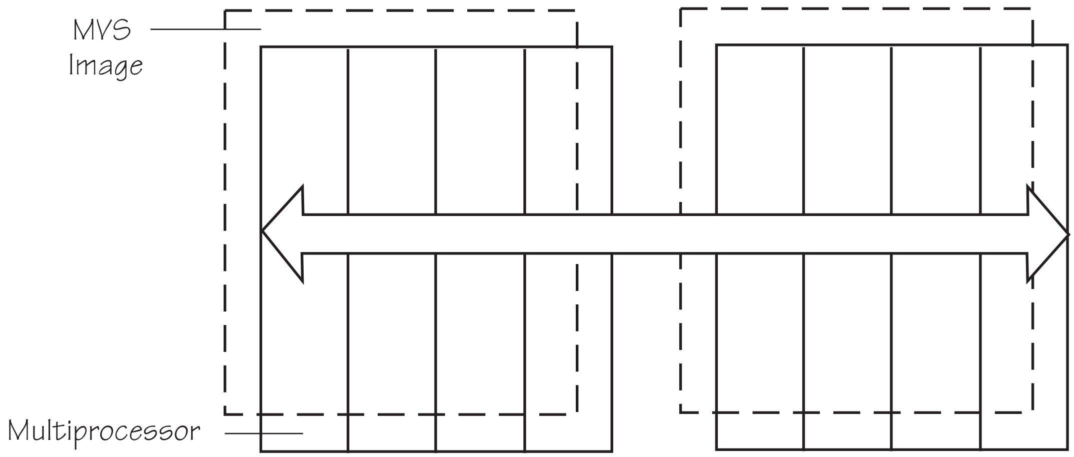

Figure 5. A Base Sysplex

Summary

Capacity
The same system capacity as loosely coupled

Availability
The same system availability as loosely coupled

Systems Management
Better and simpler systems management than loosely coupled

- Greater degree of communication and cooperation among systems
- Introduction of XCF as a common communication mechanism that provides high availability
- A more unified system image—single MVS console to manage all components

Parallel Sysplex

Since the introduction of the sysplex, IBM has developed technologies that enhance sysplex capabilities. The Parallel Sysplex supports a greater number of systems and significantly improves communication and data sharing among those systems.

High performance communication and data sharing among a large number of MVS systems could be technically difficult. But with the Parallel Sysplex, high performance data sharing through a new coupling technology (coupling facility) gives high performance multisystem data sharing capability to authorized applications, such as MVS subsystems. Use of the coupling facility by subsystems, such as Information Management System (IMS), ensures the integrity and consistency of data throughout the entire sysplex.

The capability of linking together many systems and providing multisystem data sharing makes the sysplex platform ideal for parallel processing, particularly for online transaction processing (OLTP) and decision support. More discussion about parallel processing follows in "Chapter 3. Parallel Processing in a Sysplex" on page 19.

In short, a Parallel Sysplex builds on the base sysplex capability, and allows you to increase the number of CPCs and MVS images that can directly share work. The coupling facility enables high performance, multisystem data sharing across all the systems. In addition, workloads can be dynamically balanced across systems with the help of new workload management functions.

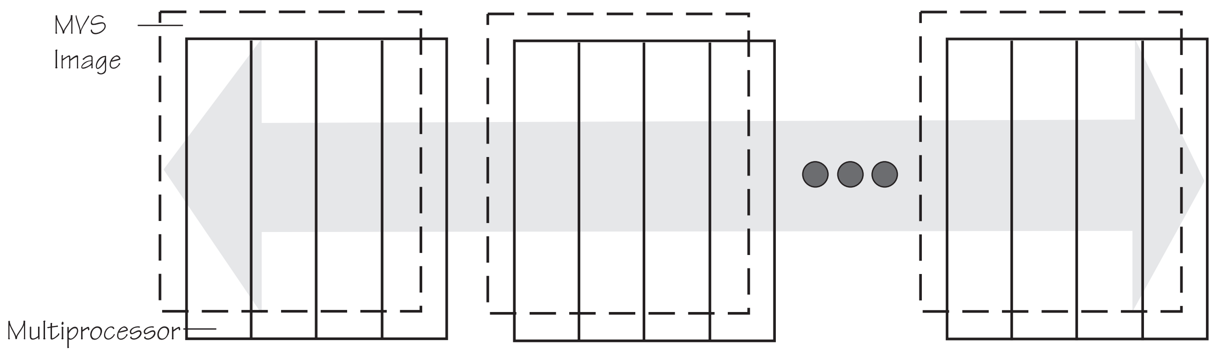

Figure 6. A Parallel Sysplex

Summary

Capacity
Increased system capacity over loosely coupled and base sysplex

Ability to add incremental capacity to match workload growth

Availability
Increased system availability over loosely coupled and base sysplex

Systems Management
Better systems management than with the base sysplex

- Multisystem data sharing capability
- Multisystem workload balancing
- Enhanced single-system image

Summary

The unique characteristics of a Parallel Sysplex can allow you to reduce your total cost of computing over prior offerings of comparable function and performance. Sysplex design characteristics mean that you can run your business continuously, even when it is growing or changing. You can dynamically add and change systems in a sysplex and configure them for no single points of failure. If your revenue depends on continuously available systems, a sysplex can protect your revenue. If you need to grow beyond the limits imposed by today's technology, a sysplex lets you go beyond those limits, and helps you avoid the complex and expensive splitting and rejoining of data centers.

The innate robustness and reliability of the MVS operating system and System/390 processors are the foundation of a sysplex. That robustness and reliability are extended to all systems in a sysplex through cross system workload balancing and data sharing using the coupling technologies. Therefore applications on multiple systems can be continuously available to end users, yet the applications are shielded behind a single-system view.

The applications that run in a sysplex are the same applications you run today. Reuse of applications and data processing skills reduce the costs of application development, re-engineering, and retraining.

Therefore, **the sysplex can reduce your overall cost of computing now and can position your data processing center for continued savings in the future**.

As You Read On...:

In this introductory chapter, you have seen how large systems have evolved into a sysplex. The sysplex itself is evolving, and the roll-out of sysplex support across software products will occur over time. For information about availability of sysplex software support, consult your IBM marketing representative.

The rest of this publication will not differentiate between the base sysplex capability and the Parallel Sysplex capability. When you see the term***sysplex***, understand it to mean the Parallel Sysplex with its enhanced capability.

Chapter 2. Data Sharing in a Sysplex

Connecting a large number of systems together brings with it special considerations, such as how the large number of systems communicate and how they cooperate to share resources. These considerations affect the overall operation of MVS systems.

A sysplex significantly changes the way MVS systems can share data. As the number of systems increase, it is essential to have an efficient way to share data across systems. The coupling facility enables centrally accessible, high performance data sharing for authorized applications, such as subsystems and MVS components, that are running in a sysplex. These subsystems and components then transparently extend the benefits of data sharing to their applications.

Use of the coupling facility substantially improves the feasibility of connecting many MVS systems together in a sysplex to process work in parallel.

This chapter describes:
- Multisystem data sharing
- The coupling facility and its exploiters

Data Sharing

Data sharing is not new. Different applications within a system often need to access the same information, sometimes to read it and other times to update it. Sometimes several copies of the data exist and with that comes the requirement of keeping all the copies identical. If the system fails, you need a way to preserve the data with the most recent changes.

Data validity is controlled by a data management system. Within a single MVS system, the data management system keeps track of which piece of data is being accessed or changed by which application in the system. It is the data management system's responsibility to capture and preserve the most recent changes to the data, in case of system failure.

When two or more MVS systems share data, each system has its own copy of a data management system. Communication between the data management systems is essential.

Therefore, multisystem data sharing hinges on high performance communication to ensure data validity among multiple data management systems and it requires high speed data accessing methods.

Data Validity

When many applications share data, within the same MVS system or among several MVS systems, they need mechanisms to guarantee data validity. Two such mechanisms are serialization and data consistency (buffer invalidation).

Serialization is a mechanism that allows control over the access and update of data. Through serialization, only one application can access and change a given piece of data at a time. As the number of applications increases, contention over a piece of data becomes more likely, and performance slows. One way to allow more applications on multiple systems to access more data at the same time is to allow serialization on smaller pieces of data. Generally, serialization at a lower level of granularity increases overhead, which slows performance. But, by using the coupling facility, data management systems can afford to serialize at a low level with acceptable performance.

Another requirement for maintaining data validity is a mechanism that ensures consistency among local copies of data. One such mechanism is buffer invalidation. This mechanism indicates to each of the data sharing systems whether their local data reflects the most recent changes.

Database managers need high speed serialization and data consistency mechanisms across systems. Not until data sharing based on the coupling facility has it been possible to efficiently share data using these mechanisms across more than two MVS systems.

Multisystem Data Access

Before sysplex, the options for accessing shared data across multiple copies of a data management system were:

- Using a single data server
- Partitioning data across multiple systems
- Sharing data between two systems

These three options are contrasted with accessing shared data using the coupling facility.

Using a Single Data Server

With the single data server option, one system controls data access and update. Other systems must funnel requests through the data server, and the cross system messaging that results can adversely affect performance. If the single data server fails, other systems can no longer access data.

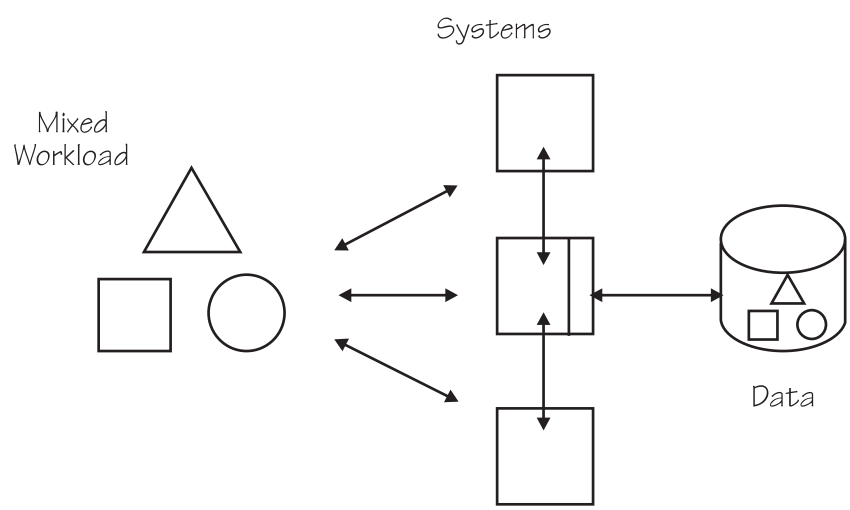

Figure 7. Data Sharing with a Single Data Server

When the data server reaches its capacity, there may be no way to expand, other than to install a larger capacity system or to go to another data sharing option.

Summary

Response Time
Limited by the capability of the single server

Throughput
Limited by the capability of the single server

Continuous Availability
Limited by the availability of the single server

Growth
Requires upgrade of server or change to a different option

Data Management
Managed as a single data management system

Partitioning Data Across Multiple Systems

Splitting the data among systems gives each data server control over a portion of data. Requests for specific data are routed to the data server that controls that data. This option gives you more capacity than a single data server, but can impose a significant systems management burden to partition the data for adequate performance.

To allow for changes in demand for data, you need to configure the size of each system to handle the peak demand for that data, which requires extra capacity. Without that extra capacity, dynamic swings in demand for a particular partition of data could make performance erratic.

If a system fails, its portion of data cannot be accessed unless there is an alternate path to the data from another system. The other system must then have enough spare capacity to assume the additional work.

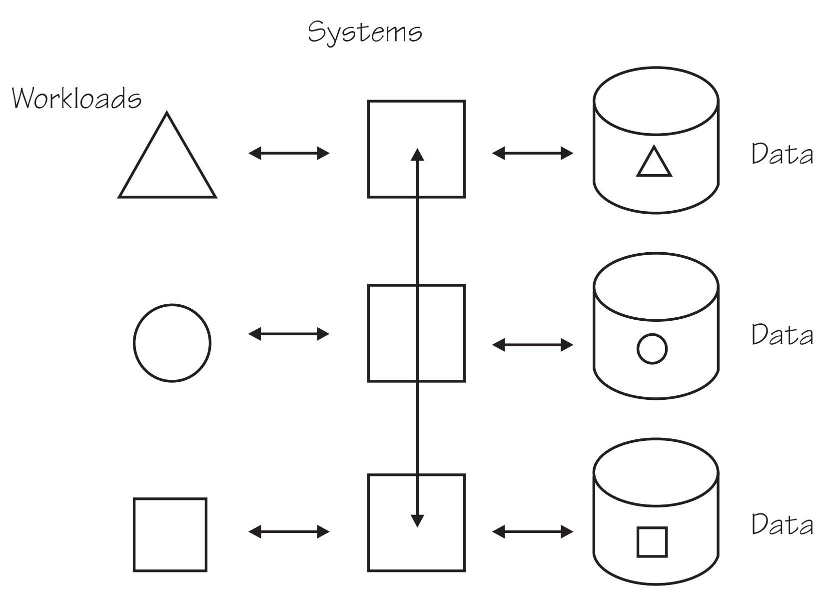

Figure 8. Data Sharing with Partitioned Data

When capacity is reached and you need to add another system, the data must be re-partitioned.

Summary

Response Time
Depends on how appropriately the data is partitioned

Throughput
Depends on the demand for data

Continuous Availability
Depends on the back-up alternate routing provisions and the capacity of the alternate system

Growth
Requires re-partitioning of the data

Data Management
Requires carefully partitioned data for adequate performance

Sharing Data Between Two Systems

Using the Information Management System Database Manager (IMS DB) without the coupling facility, you can efficiently share data between *two* different MVS systems. Whenever one of the sharing IMS DB systems needs to access the database, it sends a message to the other.

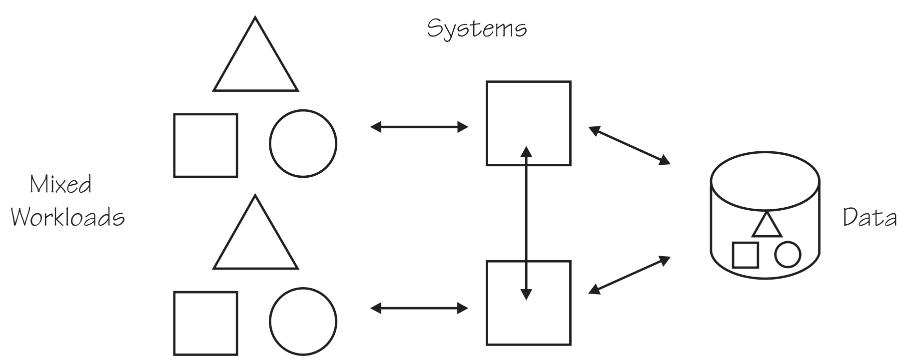

Figure 9. Data Sharing Between Two Systems

If you increase the number of IMS DB systems sharing the database to more than two, the number of messages from one sharing IMS system to the others would slow all the systems down.

Summary

Response Time
Acceptable

Throughput
Limited by the two IMS DB systems

Continuous Availability
Limited by the two IMS DB systems

Growth
Limited by the two MVS systems

Data Management
Managed as a single data management system

Sharing Data Using the Coupling Facility Technology

In contrast to the options discussed, data sharing based on the coupling facility makes it practical for you to have read/write data sharing among more than two MVS systems. The coupling facility allows data management systems to communicate so that they can directly share data. There is no single system creating a bottleneck; the data does not need to be partitioned or re-partitioned when you add another system; and there is no longer a two-system limitation.

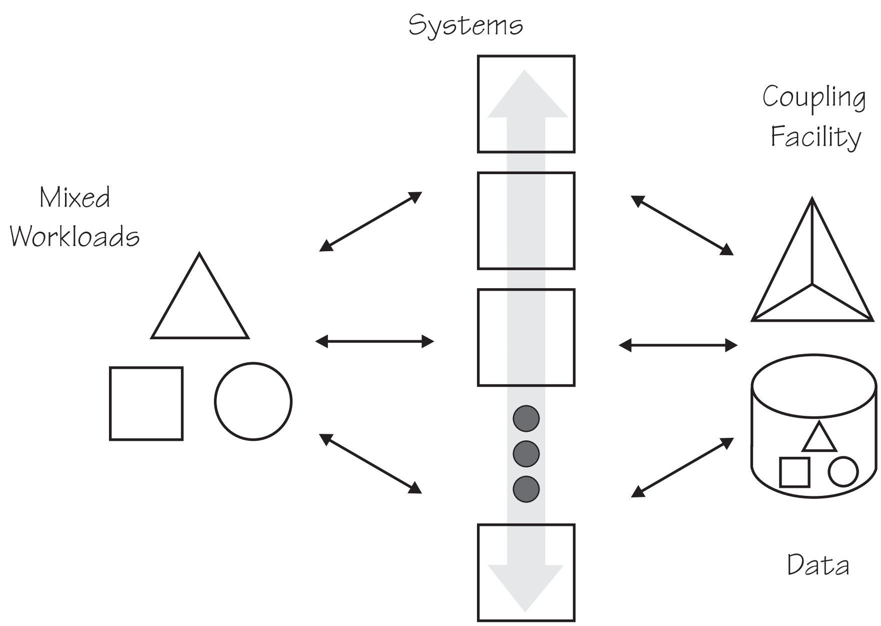

Figure 10. Sysplex Data Sharing Using the coupling facility

Summary

Response Time
Relieves poor response time due to system overload

Throughput
Increases throughput far beyond the current limits

Continuous Availability
Extends availability far beyond the current limits

Growth
Incremental expansion is not disruptive to existing systems

Data Management
Managed as a single data management system

The Coupling Facility Technology

The technology that makes high performance sysplex data sharing possible is a combination of hardware and software services available in the supporting MVS release.

A coupling facility includes the following:

- S/390 9674 Coupling Facility (standalone models)
- Logical partition of a S/390 9672 Parallel Sysplex Transaction Server or 9672 Parallel Sysplex Enterprise Server

High bandwidth fiber optic links called coupling facility channels, provide high speed connectivity between the coupling facility and systems directly connected to it.

Within the coupling facility, storage is dynamically partitioned into ***structures***. MVS services manipulate data within the structures. Each of the following structures has a unique function:

- Cache structure - Supplies a mechanism called buffer invalidation to ensure consistency of cached data. The cache structure can also be used as a high-speed buffer for storing shared data with common read/write access.
- List structure - Enables authorized applications to share data that is organized in a set of lists, for implementing function such as shared work queues and shared status information.
- Lock structure - Supplies shared and exclusive locking capability for serialization of shared resources down to a very small unit of data.

Exploiters of the Coupling Facility

Authorized applications, such as subsystems and MVS components in the sysplex, can use the coupling facility services to cache data, share queues and status, and access sysplex lock structures in order to implement high performance data sharing and rapid recovery from failures. The subsystems and components transparently provide the data sharing and recovery benefits to their applications.

Some IBM data management systems that are using the coupling facility include database managers and a data access method.

Information Management System Database Manager (IMS DB): IMS DB is IBM's strategic hierarchical database manager. It is used for numerous applications that depend on its high performance, availability, and reliability. A hierarchical database has data organized in the form of a hierarchy (pyramid). Data at each level of the hierarchy is related to, and in some way dependent upon, data at the higher level of the hierarchy.

IMS database managers on different MVS systems can access data at the same time. By using the coupling facility in a sysplex, IMS DB can efficiently provide data sharing for more than two MVS systems and thereby extends the benefits of IMS DB data sharing. IMS DB uses the coupling facility to centrally keep track of when shared data is changed. IRLM is still used to manage data locking, but does not notify each IMS DB of every change. IMS DB does not need to know about changed data until it is ready to use that data.

DATABASE 2 (DB2): DB2 is IBM's strategic relational database manager. A relational database has the data organized in tables with rows and columns.

DB2 data sharing support allows multiple DB2 subsystems within a sysplex to concurrently access and update shared databases. DB2 data sharing uses the coupling facility to efficiently lock, to ensure consistency, and to buffer shared data. Similar to IMS, DB2 serializes data access across the sysplex through locking. DB2 uses coupling facility cache structures to manage the consistency of the shared data. DB2 cache structures are also used to buffer shared data within a sysplex for improved sysplex efficiency.

Virtual Storage Access Method (VSAM): VSAM, a component of DFSMS, is an access method rather than a database manager. It is an access method that gives CICS and other application programs access to data stored in VSAM-managed data sets.

DFSMS supports a new VSAM data set accessing mode called record level sharing (RLS). RLS uses the coupling facility to provide sysplex-wide data sharing for CICS and the other applications that use the new accessing mode. By controlling access to data at the record level, VSAM enables CICS application programs running in different CICS address spaces, called CICS regions, and in different MVS images, to share VSAM data with complete integrity. The coupling facility provides the high performance data sharing capability necessary to handle the requests from multiple CICS regions.

Other Exploiters of the Coupling Facility

In addition to data management systems, there are other exploiters of the coupling facility, such as SecureWay Security Server (RACF), and JES2. Transaction management systems also exploit the coupling facility to enhance parallelism.

Chapter 3. Parallel Processing in a Sysplex

Parallel processing is a technology that increases application availability. It gives you more servers for enhanced throughput; and if the enhanced throughput relieves system overload, parallel processing will visibly reduce end-user response time.

One reason for setting up a sysplex is to enable parallel processing. Special purpose parallel systems for engineering and scientific applications have been available for years. The sysplex, however, can provide parallel processing for commercial online transactions.

The sysplex brings together parallel capability on many levels—in the applications, the supporting sysplex software, and the hardware. Applications can be replicated to run in parallel on a single tightly coupled multiprocessing system. When you run parallel applications in a sysplex with subsystems that support enhanced communication across systems, you have a greater degree of parallelism. And when you add hardware to that greater degree of parallelism, hardware with the capability for modular, incremental growth—you have the broad based parallelism that is available in a sysplex.

This chapter discusses:

- A definition of parallel processing
- Running work in parallel
- Software that enables parallelism for OLTP
- Hardware that enables parallelism

What is Parallel Processing?

Parallel processing in a sysplex is the ability to simultaneously process a particular workload on multiple CPCs, each of which may have multiple CPs. A CP can generally process one unit of work at a time. Typically the more CPs in a system, as in a tightly coupled multiprocessor, the more work can be processed simultaneously. When you have multiple CPCs processing the same workload, you can significantly reduce elapsed processing time.

Parallel processing is a type of multiprocessing. For the purposes of this publication, the distinction between multiprocessing and parallel processing is in whether the same kind of work is being processed. While multiprocessing is the simultaneous processing of work, it is not necessarily the same kind of work; parallel processing in a sysplex is the simultaneous processing of one type of work, whether it is many small but distinct units of work, or a larger unit of work that is broken down to run in parallel.

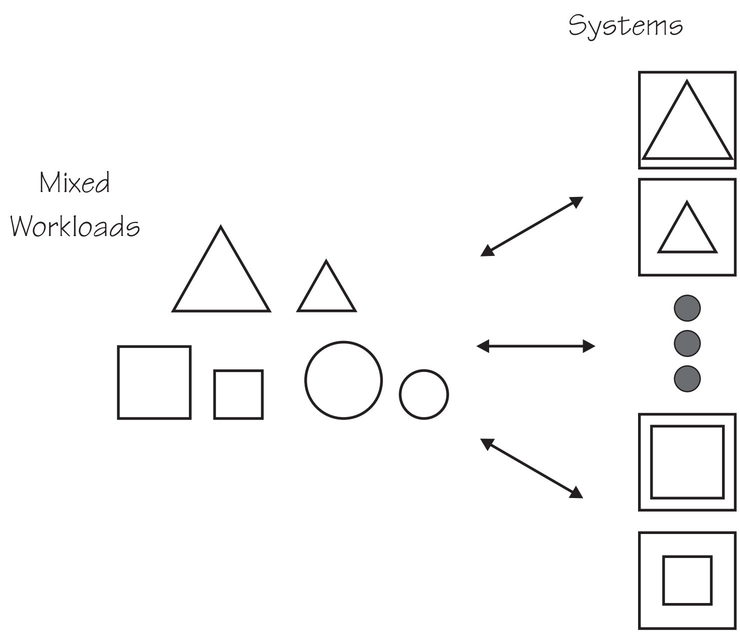

Figure 11. Multiprocessing

In a multiprocessing environment, at a particular moment in time, one CP might be running a payroll batch application, while another CP processes online transactions. The different work runs at the same time. If the payroll batch application takes hours to complete, and if the application is structured as a single task, it can only use the power of one CP. The fact that there is more than one CP available does not split up the long-running batch application. But if the batch application is structured as multiple tasks, it can be processed simultaneously on many CPs—in other words, in parallel.

Similarly, if there are more online transactions than one CP can handle, adding more CPs allows the online transactions to run in parallel.

Therefore, parallel processing is the ability to use many CPs to simultaneously process a particular workload; and by running a workload in parallel, it is possible to make a noticeable improvement in the response time for the end user.

In a sysplex, it is possible to have one workload running on many tightly coupled multiprocessors that are sharing data using the coupling facility. How many CPs you have depends on the configuration. For example, a configuration of 16 CPCs with 6 CPs per system gives you a total of 96 CPs that can process a workload in parallel.

Problems Solved by Parallel Processing

Parallel processing in a sysplex can be used for two different programming problem situations: when the number of transactions noticeably slows response time and when a long-running application takes too long to complete. Both situations can adversely affect end-user response time.

Increased Number of Transactions: When the number of incoming transactions exceeds the processing capability, transactions are delayed and response time increases. The following graphic shows how a large number of transactions can build up when they exceed the processing capability.

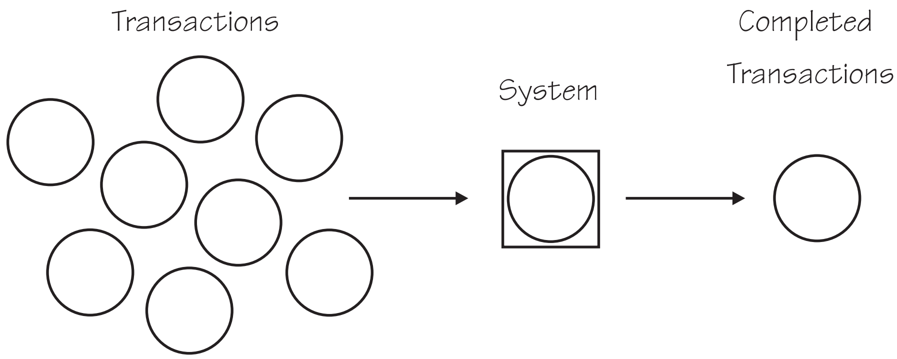

Figure 12. Workload that Exceeds Processing Capability

By adding processors and running transactions in parallel, you can reduce end-user response time by increasing throughput. The following graphic shows how adding processing capability increases throughput.

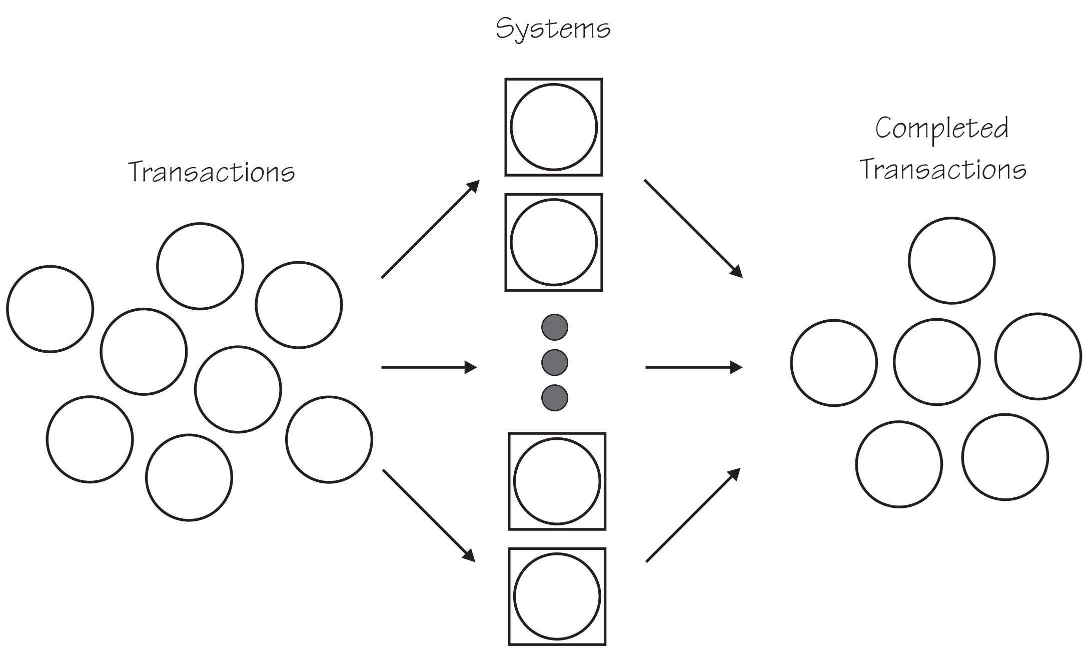

Figure 13. Increasing Throughput for a Workload

Similarly in your environment, when an increased number of transactions slows throughput, you can use a sysplex and the coupling facility to increase the number of CPs running in parallel.

Lengthy Application: To shorten the processing time of a single, long-running application, subdivide it into smaller units of work, if possible. These smaller units can then run in parallel. The following graphic shows how a lengthy application is processed one piece at a time when there is only one CP.

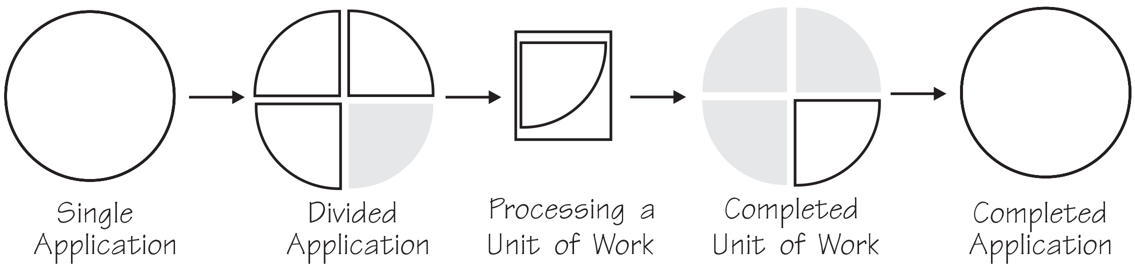

Figure 14. Processing a Lengthy Application

When the lengthy application is divided into smaller units of work and there are many parallel processors, the smaller units of work can run simultaneously and in parallel. The following graphic shows how a lengthy application can be processed in parallel by dividing it among many parallel processors.

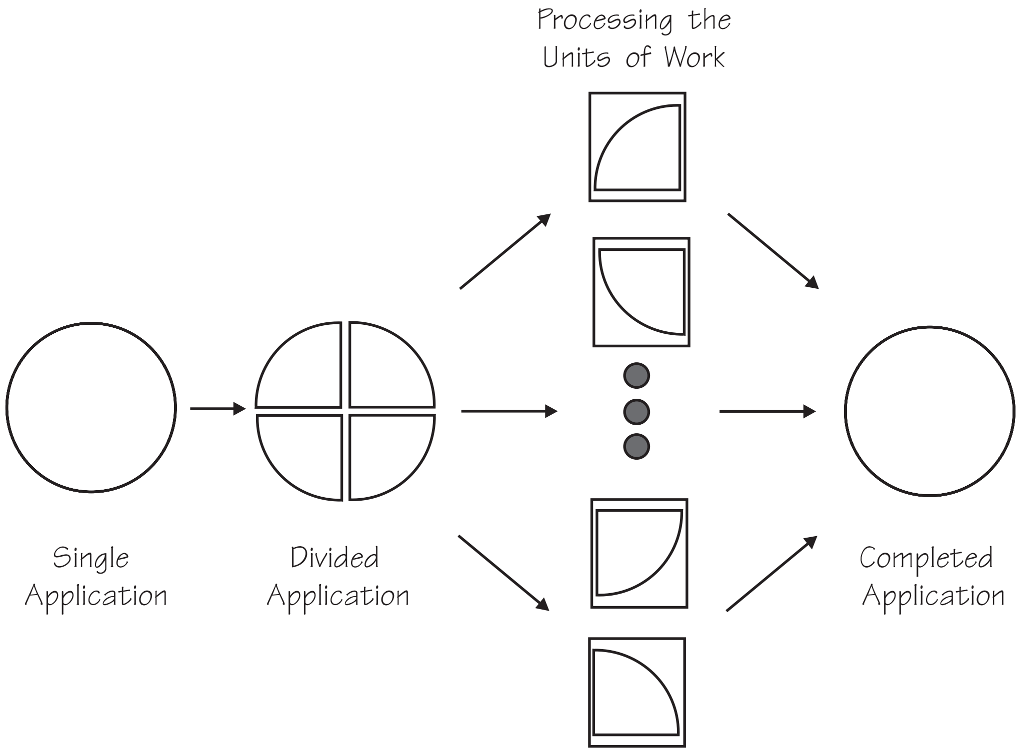

Figure 15. Decreasing the Processing Time of a Lengthy Application

You can use this same approach in a sysplex to speed up the processing of a long-running application, such as a complex query.

Running Work in Parallel

MVS is known for its strength and dependability in processing applications that solve large business problems. These are some characteristics of a business problem, each of which makes it a *large* problem:

- Volume of work, up to millions of transactions per day
- Number of end users, up to thousands
- Amount of storage required, such as terabytes of data
- Bandwidth requirements for sending data in a network, such as megabytes of data in seconds
- Stringent requirements for continuous availability (24 hours a day, 7 days a week, 365 days a year)

Some common MVS applications that deal with large business problems and can be made to run in parallel are batch, query, and online transaction processing (OLTP).

A batch workload consists of multiple job streams, which can be unrelated and can therefore run in parallel. A single job stream, however, runs in sequence and usually cannot take advantage of parallelism. When a single job stream is lengthy, it can take considerable time to complete. IBM provides a product called Parallel Sysplex, to help you exploit parallelism in a single job stream within a single MVS image.

Queries are a diverse type of workload and usually rely on partitioned data when running in parallel. To help you process queries in parallel, IBM is offering a parallel query server. The parallel query server is a self-contained system made up of hardware, software, and database services. The system is fully installed, customized, and serviced by IBM. The parallel query server is used to process large queries in a shared relational database. For more information about the parallel query server, see your IBM marketing representative.

OLTP applications are a common MVS application type that can take advantage of parallelism. They are the focus of the following topics, because there is extensive software support for OLTP processing in a sysplex, and the number of OLTP applications are on the increase.

Online Transaction Processing (OLTP)

OLTP applications are used by banks, airlines, insurance companies, and other businesses that give online users direct access to information. The OLTP applications process units of work, called transactions. A single transaction might request a bank balance; another might update that balance to reflect a deposit.

An application can be replicated to run in parallel on a tightly coupled multiprocessor so that it can simultaneously process multiple transactions. The degree of parallelism depends on the system software, the transaction manager, and the hardware that the OLTP application runs on.

Software that Enables Parallelism for OLTP

OLTP applications use system software services provided by transaction managers and they access data controlled by data management systems. To balance the transactions dynamically across systems, MVS workload manager, a component of MVS, works with the transaction managers. This section first discusses the transaction managers and then describes workload management.

Transaction Managers

From the time the transaction manager obtains the transaction from the network until it returns the completed transaction to the network, it is responsible for:

- Sending the transaction to the appropriate application
- Converting transaction data to a form that can be used
- Reconverting results to a form that can be displayed on an output device

If the transaction manager can simultaneously process copies of an application in multiple address spaces within a single system, an OLTP application can run in parallel.

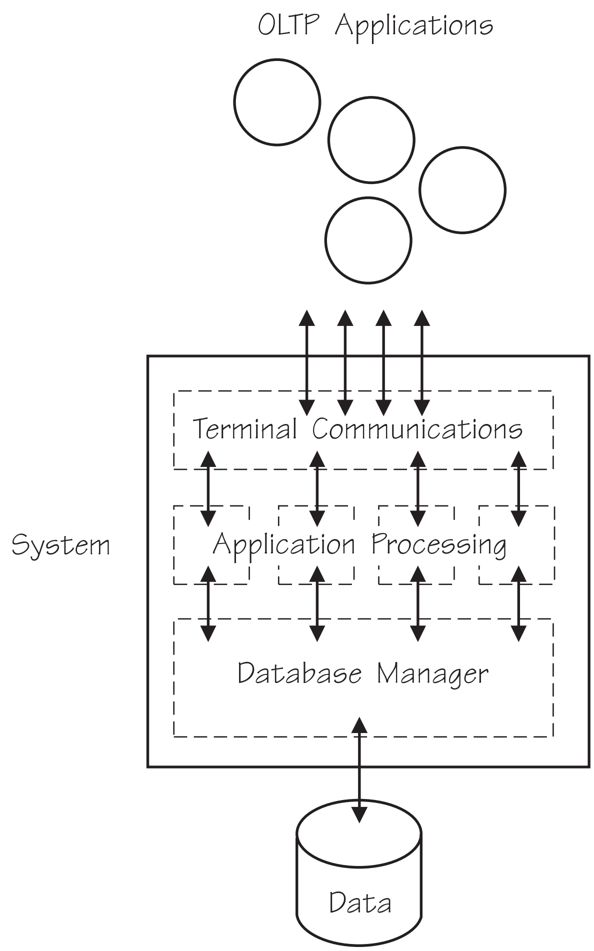

Figure 16. Parallel OLTP Applications in a Single System

If the transaction manager can take advantage of sysplex support and simultaneously process copies of an application across different MVS systems, the OLTP application can run with a greater degree of parallelism.

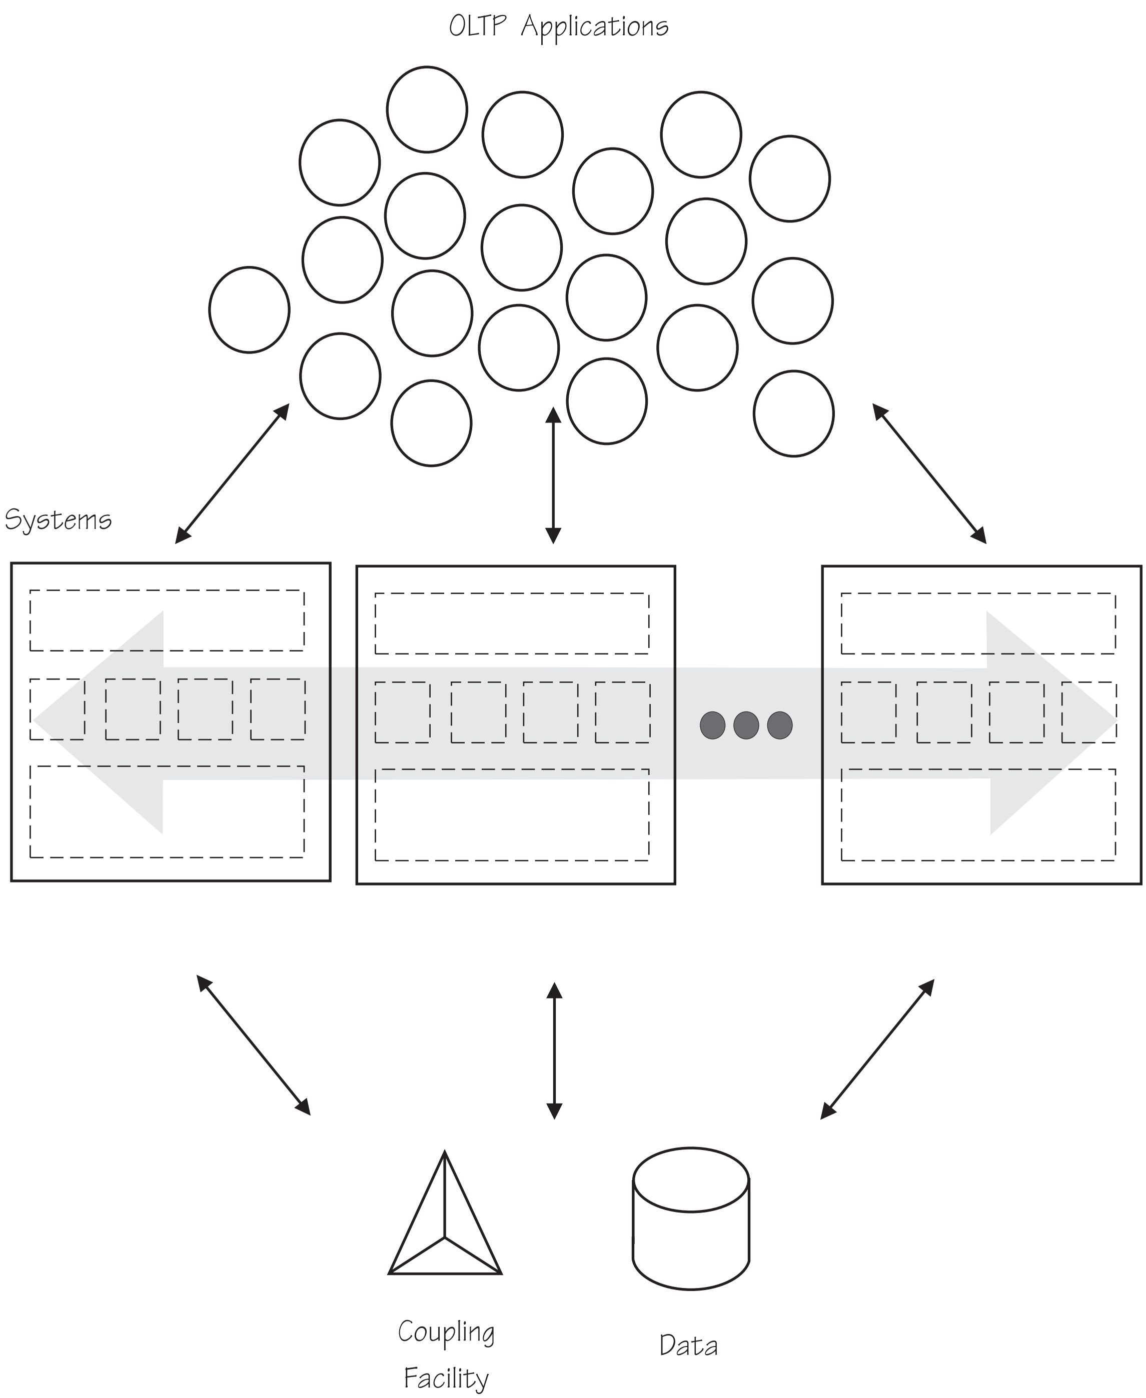

Figure 17. Parallel OLTP Applications in a Sysplex Parallel Environment

Two MVS transaction managers, which have provided parallel support in a single system, provide an additional degree of parallel support in a sysplex. The two transaction managers are Customer Information Control System (CICS) and Information Management System Transaction Manager (IMS TM).

Customer Information Control System (CICS): CICS enables online transactions entered at terminals to be processed concurrently by user-written application programs. Originally CICS ran in a single MVS address space that was responsible for terminals, applications, and files.

Using CICS multi-region operation, you can separate function into individual regions, namely terminal-owning regions (TORs), application-owning regions (AORs), and file-owning regions (FORs). These individual resource-owning regions can then be linked together and managed in what is called a CICSplex.

One of the benefits of separating the CICS functions into a number of resource-owning regions is to take advantage of multiprocessor systems, as each region (address space) can run on a different CP. You can structure a CICS environment such that the application-owning regions can process transactions in parallel.

Replicating regions and running multiple regions on multiple systems in a sysplex further increases the parallelism. The greater the number of regions, the greater the CICS availability in the event of a system failure.

Information Management System Transaction Manager (IMS TM): The IMS Transaction Manager (IMS TM) is an IBM transaction processing product that is designed to take advantage of multiprocessors. IMS TM is a multi-address space product so that one IMS TM running under one MVS can effectively use tightly coupled multiprocessors.

IMS TM can also take advantage of the sysplex. Several IMS TMs can coordinate work among themselves by using IMS's Multiple Systems Coupling (MSC) support. Through MSC, transactions entered from one IMS can be routed to another IMS for processing and the results returned to the original IMS. To the end user, the transfer of information from one IMS to another is transparent.

Transaction message traffic workloads can be balanced using MSC. IMS will continue to enhance and unify management of IMS systems and present a single-system view of IMS in a sysplex.

Data Management Systems

When an OLTP application needs to access data, the data management system controls the data access and update. Data management systems are responsible for validity of the data, providing a way to recover data in case of failure, and a way to secure data from unauthorized access.

If the data management system takes advantage of the coupling facility, the OLTP applications can avoid shipping data back and forth between systems. By using the coupling facility for data sharing and for queuing and status control, data management systems and transaction managers can perform their duties more quickly and therefore manage more systems efficiently. This increase in the number of systems that can be managed, increases the degree of parallelism even further.

The data management systems that take advantage of the coupling facility are IMS DB, DB2, and VSAM.

Managing Work in a Sysplex

Along with parallelism comes the need for simple and dynamic multisystem workload balancing. When you are in an environment with multiple systems, the set of performance issues changes. Existing mechanisms for managing system performance are complex and single-system oriented.

To reduce the complexity of managing a sysplex, MVS workload management provides dynamic sysplex-wide management of system resources. MVS workload management is the combined cooperation of various subsystems (such as CICS, IMS, and VTAM) with the MVS workload manager (WLM) component. An installation defines performance goals and a business importance to workloads through WLM. Workload management focuses on attaining these goals through dynamic resource distribution.

This type of workload management is different from the way workloads were previously managed. The new emphasis is on defining performance goals for work, and having MVS and the subsystems adapt to meet the goals.

Defining Performance Goals

An installation defines performance goals in a ***service policy***. Service policies are defined through an ISPF application and they set goals for all types of MVS-managed work. An installation can create multiple service policies to adjust performance goals for different periods of time.

The scope of a service policy is the sysplex. Each service policy has a name, and can be activated by the ISPF application, an operator command, or an automation package. Only one policy can be active at a time. When it is activated, all systems in the sysplex process towards the goals defined in the policy.

Workload management coordinates and shares performance information across the sysplex. Each MVS system handles its own system resource management and dynamically matches resources to work according to the goals defined in the service policy. During processing, the system monitors how well the goals are being met and adapts accordingly as the environment changes. If there is contention for resources, each system makes the appropriate trade-offs based on the importance of the work and how well the goals are being met. This way, all systems can cooperate to protect work that is critical to your installation.

Reporting

Resource Measurement Facility (RMF) combines MVS system management facilities (SMF) data for the sysplex, and reports how well the sysplex is doing to achieve the goals defined in the service policy. In addition, execution delay information is available in SMF records that show where delays are occurring. If there is a problem, you can use this information to help adjust the performance goals, focus on specific subsystems having a problem, or make work scheduling adjustments.

Hardware that Enables Parallel Processing

Sysplex parallel processing depends on coupling facility technology and the ability to connect to it. Several models of S/390 or Z900 processors are able to connect to the coupling facility including large capacity water-cooled processors and smaller capacity air-cooled processors. Coupling facilities include the S/390 9674 Coupling Facilities. In addition, processors that make use of complementary metal oxide (CMOS) technology (S/390 9672 Parallel Sysplex Transaction Servers and 9672 Parallel Sysplex Enterprise Servers) can have one or more logical partitions configured as integrated coupling facilities while other logical partitions can run as parallel transaction servers to process work.

The significance of adding CMOS processors to a sysplex is the improved price/performance and incremental growth possible by using large numbers of relatively small CPCs to process work in parallel. But, regardless of processor size, the use of new hardware technology makes sysplex parallel processing a cost effective alternative.

Chapter 4. Products in a Sysplex

A sysplex is not a single product that you install in your data center. Rather, a sysplex is a collection of products, both hardware and software, that work together.

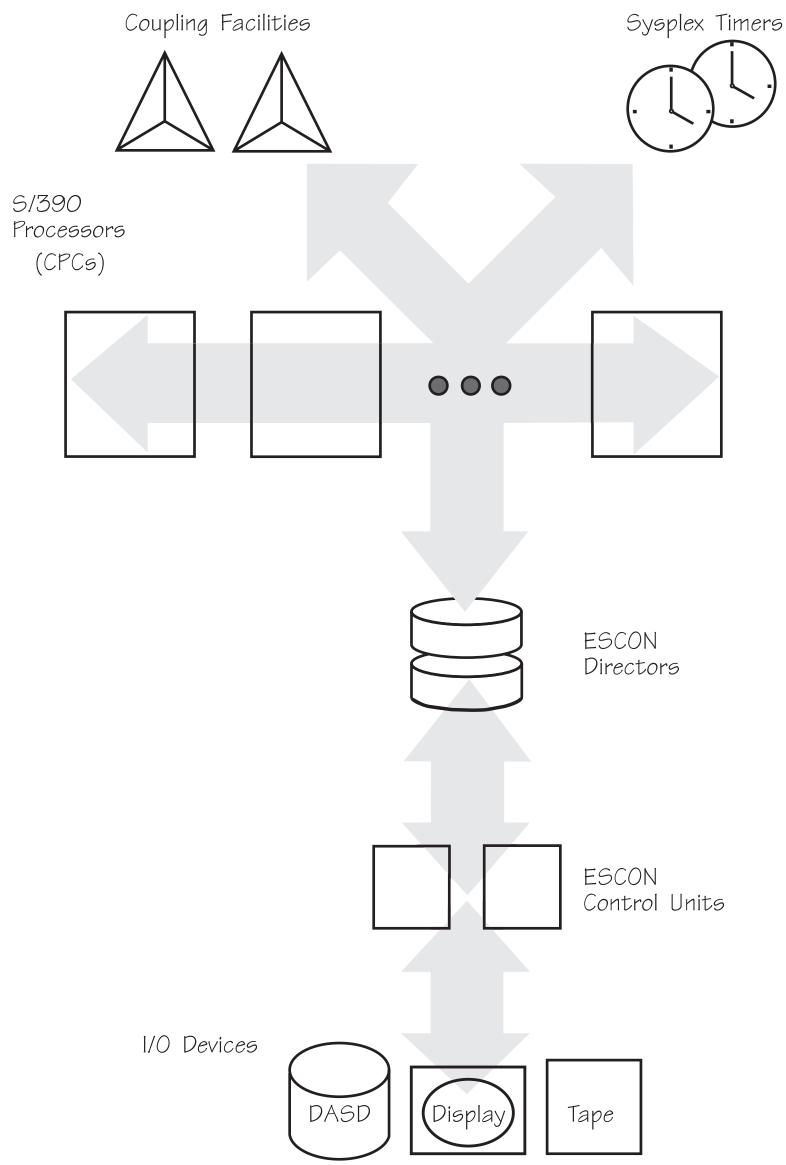

Figure 18. Products in a Sysplex

The products that work together in a sysplex can be categorized as hardware, system software, networking software, data management, transaction management, and systems management. Following are sysplex products grouped by area.

Note: A sysplex is not limited to the products described in this chapter. Additional IBM and non-IBM products might be able to run in a sysplex or exploit sysplex function. For a current list of products that can take advantage of a sysplex, see your IBM marketing representative.

Hardware

Hardware for the Parallel Sysplex includes the following.

Coupling Facility

The coupling facility is a CMOS processor (a standalone S/390 9674 Coupling Facility or a logical partition defined in the 9672 Parallel Transaction Server or the 9672 Parallel Enterprise Server). It allows high performance multisystem data sharing among certain S/390 or z900 processors.

MVS services allow authorized applications, such as subsystems and MVS components, to use the coupling facility to cache data, exchange status, and access sysplex lock structures in order to implement high performance data sharing and rapid recovery from failures.

Coupling Facility Channels

Coupling facility channels are high bandwidth fiber optic links that provide high speed connectivity between the coupling facility and the CPCs that use the coupling facility. Coupling facility channels are directly attached between the CPCs and the coupling facility.

Processors

Note: All ES/9000 processors can participate in a base sysplex, but only the following processors can link to the coupling facility.

- 9021 711-based models

  The 9021 711-based models are water-cooled CPCs, which were announced in February, 1993.
- 9121 511-based models

  The 9121 511-based models are air-cooled CPCs, which were announced in February, 1993.
- S/390 9672 Parallel Transaction Server and 9672 Parallel Enterprise Server

  A S/390 9672 Parallel Transaction Server and 9672 Parallel Enterprise Server comes with an integrated coupling facility and flexible packaging options, and it runs as a parallel transaction server. In addition, an easy-to-use console provides cross system management capability. For a description of the console, see “Hardware Management Console” on page 35.
- z900 processor models

Sysplex Timer

The Sysplex Timer is a unit that synchronizes the time-of-day (TOD) clocks in multiple CPCs in a sysplex. The time stamp from the Sysplex Timer is a way to monitor and sequence events within the sysplex.

Enterprise Systems Connection (ESCON)

ESCON channels have a channel-to-control-unit I/O interface that uses optical cables as a transmission medium. ESCON Directors (ESCDs) add dynamic switching capability for ESCON channels, further increasing connectivity and device sharing.

ESCON I/O and the interconnect technologies become very important in a sysplex to speed access to shared data on disk or tape, and to enhance communication among systems. In addition, they offer improved availability.

ESCON Control Units and I/O Devices

To ease the management of logical paths and increase connectivity, use control units and devices that provide a higher number of logical paths, such as:

- 3990 Storage Control Model 6
- 9343 Storage Controller and 9345 Direct Access Storage Device
- I/O devices that are ESCON capable (3490E, 3172, 3174, 3745, 3746)

System Software

System software for the Parallel Sysplex includes the following.

OS/390 or z/OS

The OS/390 or z/OS operating systems are known for their ability to handle thousands of concurrent interactive users and for its sophisticated data management, reliability, security, and auditing features in a commercial environment.

MVS, an element of the operating systems, enables the new coupling technology, extends system support beyond the previous limit of 8 MVS systems, and increases the number of I/O devices allowed per MVS system beyond the previous limit of 4096. It is the platform for simplified systems management of a sysplex, including configuration management, availability management, workload management, and single-image operations.

JES2 or JES3

These two job entry subsystem products control job queues and dispatch work in a sysplex. JES2 expands its 7-member multi-access spool (MAS) capability to a 32-member MAS. Additionally, JES2 uses the coupling facility to enhance its performance. JES3 can support a maximum of 32 members in a JES3 complex and supports ESCON CTCs or the coupling facility.

DFSMS

DFSMS enables the automatic placement, migration, backup, recall, recovery, and deletion of data for MVS. An SMS configuration can contain 32 names, but the names can be a combination of system names and system group names. DFSMS also supports the increased number of I/O devices beyond the previous limit of 4096.

Networking Software

Networking software for the Parallel Sysplex includes the following.

Virtual Telecommunications Access Method (VTAM)

VTAM monitors and controls the activation and connection of resources in a network. When VTAM is part of the sysplex, duplicate applications on different systems in a VTAM network can identify themselves by a single generic resource name. By distributing sessions among a number of duplicate resources under a single name, instead of to a uniquely named single resource, VTAM is capable of balancing the session workload.

Data Management

Data management for the Parallel Sysplex includes the following.

Information Management System Database Manager (IMS DB)

IMS DB is a database manager of hierarchical data that provides multisystem support across two MVS systems. By using the coupling facility, IMS DB extends data sharing across multiple MVS systems.

DATABASE 2 (DB2)

DB2 is a database manager of relational data. DB2 provides data sharing support that uses the coupling facility for read/write sharing of DB2 databases within a sysplex.

Virtual Storage Access Method (VSAM) - component of DFSMS

VSAM is an access method that gives application programs access to data stored in VSAM-managed data sets. DFSMS supports a new VSAM data set accessing mode called record level sharing (RLS). VSAM RLS provides sysplex data sharing and will use the coupling facility for locking and data caching.

Transaction Management

Transaction management for the Parallel Sysplex includes the following.

Customer Information Control System (CICS/ESA and CICS Transaction Server)

CICS/ESA and CICS Transaction Server are transaction managers that provide services for online transactions. Both CICS/ESA and CICS Transaction Server enhance the multiregion operation (MRO) links to provide more efficient cross system communication for CICS. They support the MVS workload manager component and VTAM generic resource as described previously. They also support sysplex data sharing for IMS DB and DB2. CICS Transaction Server provides support for record-level sharing of data for VSAM.

Information Management System Transaction Manager (IMS TM)

IMS TM is a transaction manager that provides services for online transactions in a multisystem environment. It cooperates with the MVS workload manager component for better performance.

Systems Management

Systems management software includes the following.

Accounting

System Management Facilities (SMF) - component of MVS

This component of MVS keeps records of resources used in the system. SMF record keeping expands beyond the previous 8 system limit. Also, SMF provides new records specifically for a sysplex.

Workload

CICSPlex System Manager/ESA (CICSPlex SM)

This product provides a single-system image for a CICSplex across a range of system management applications. These applications include dynamic workload balancing and separation in cooperation with the MVS workload manager, CICS master terminal operations, online access to CICS data and statistics, and the ability to detect exceptional events based on CICS state data for an operator or for automation purposes.

Operations Planning and Control (OPC/ESA)

OPC/ESA is a product that helps you manage your installation's batch production workload. From one system you can plan, control, and monitor all workloads. OPC/ESA has been enhanced to present a single-system image in a sysplex environment. This enhancement will give a TSO user access to the OPC/ESA dialog even when the OPC/ESA Controller and the TSO user are active on different MVS systems.

Workload Manager (WLM) - component of MVS

To manage workloads throughout a sysplex, WLM cooperates with transaction and resource managers that span systems. WLM uses customer-defined policies for performance objectives to help balance workloads.

Operations

System Automation for OS/390

System Automation for OS/390 provides system automation for subsystem start and shutdown, automation for messages, timers, system data set offload, and policy dialogs. System Automation for OS/390 also provides enterprise resource monitoring through a workstation graphical interface where the enterprise can include one or more sysplexes with coupling facilities. Where appropriate, resources are monitored at the sysplex level.

System Automation for OS/390 also eases the task of propagating automation policy to multiple systems. System Automation for OS/390 consolidates access to other products (for example, SDSF and RMF) from a single workstation.

System Automation for OS/390 also extends the control, monitoring, and automation capabilities of the NetView program to provide consolidated management for hardware (system) and software (operator) consoles in heterogeneous System/390 environments. System Automation for OS/390 supports both automated and manual operations on a sysplex.

Hardware Management Console

An enhanced hardware system console provides a single point of control for the systems in a S/390 9672 Parallel Transaction Server and 9672 Parallel Enterprise Server. Using icons and a mouse, you can operate one system, and you can perform an operation for all systems or for a user-defined group of systems.

System Display and Search Facility (SDSF)

SDSF helps you efficiently monitor and control the operation of an MVS/JES2 system. It supports the increased number of systems that can be defined in a JES2 multi-access spool and the MVS workload manager component.

Performance

Resource Measurement Facility (RMF)

RMF gathers performance data about system resources and writes SMF records 70-79. To reduce the complexity of performance management in a sysplex, RMF introduces the concept of a single system view for reporting on sysplex-wide data. It also supports the coupling facility and the MVS workload manager component.

Enterprise Performance Data Manager/MVS (EPDM)

EPDM is a system log data set post processor and reporter that reads any sequential log data set. It uses DB2 to support data summarization and reporting on new and changed data for sysplex-related SMF records. Reporting can be done online through an OS/2 dialog or through ISPF, and also in batch.

Service Level Reporter (SLR)

SLR is a system log data set post processor and reporter that reads any sequential log data set. It supports data summarization and reporting on new and changed data for sysplex-related SMF records. Reporting can be done online through ISPF and also in batch.

Security

SecureWay Security Server (RACF)

The Security Server element of OS/390 or z/OS provides security for the MVS system. The Security Server operates in a sysplex environment and takes advantage of the coupling facility to reduce I/O to the RACF database. It also provides additional systems management function.

Configuration

System Automation for OS/390

System Automation for OS/390 helps manage the complexities of active I/O configurations. When running on a sysplex, System Automation for OS/390 manages the sysplex as a unified system image. Using System Automation for OS/390, you can graphically view and change the connectivity between systems and I/O objects (shared and non-shared, ESCON and non-ESCON).

Hardware Configuration Definition (HCD)

HCD is the configuration management program for MVS. It replaces the MVS configuration program (MVSCP) and creates input statements to the I/O configuration program. For S/390 9672 Parallel Transaction Server and 9672 Parallel Enterprise Server models, HCD is the focal point for all the I/O configuration data set (IOCDS) management function. HCD can also define I/O for the coupling facility including the channel paths associated with the coupling facility channels.

Chapter 5. Introduction to Sample Sysplex Configurations

In a sysplex, you can think of MVS systems and system hardware components as building blocks that you can select for specific purposes and replicate or combine in different configurations.

Following are some sample sysplex configurations. All configurations assume:
- Shared DASD with ESCON connections
- Appropriate levels of enabling software
- Possible connection to a network

Large Systems Data Sharing Sysplex

This configuration connects two large CPCs in a sysplex that uses the coupling facility to share data. The CPCs can be either 9021 711-based models, 9121 511-based models, or one of each.

Benefits: This configuration combines the inherent strengths of large systems with multisystem data sharing provided by the coupling facility. It can be used for both:
- An easy migration path for initiating and testing data sharing using the coupling facility
- A reliable production environment

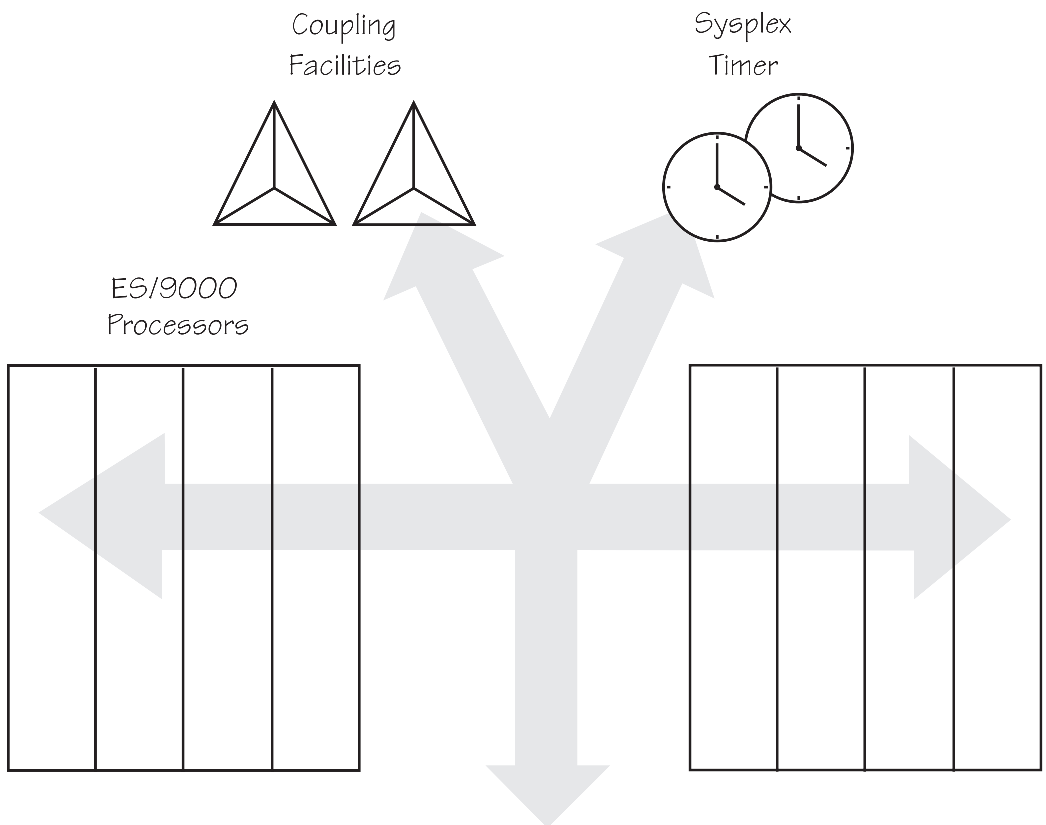

Figure 19. Large Systems Sysplex Configuration

The parts of this configuration are:

- ES/9000 large processor, such as:
  – 9021 711-based CPC
  – 9121 511-based CPC
- A standalone coupling facility (such as the 9674 Coupling Facility, with an optional second coupling facility as backup)
- A Sysplex Timer with the high availability feature

Parallel Sysplex with the 9672 Parallel Transaction Server or 9672 Parallel Enterprise Server

This separately orderable sysplex configuration consists of an independent S/390 CMOS processor (9672 Parallel Transaction Server or 9672 Parallel Enterprise Server) with an integrated coupling facility defined in a logical partition.

Benefits: In addition to the standard sysplex benefits, such as improved price/performance, greater parallel capacity for a single workload, and environmental savings (space, power, cooling, and weight), this configuration:

- Offers the reliability of a production environment
- Allows non-disruptive incremental growth, a single CPC at a time
- Can be easily added to an existing sysplex to increase capacity

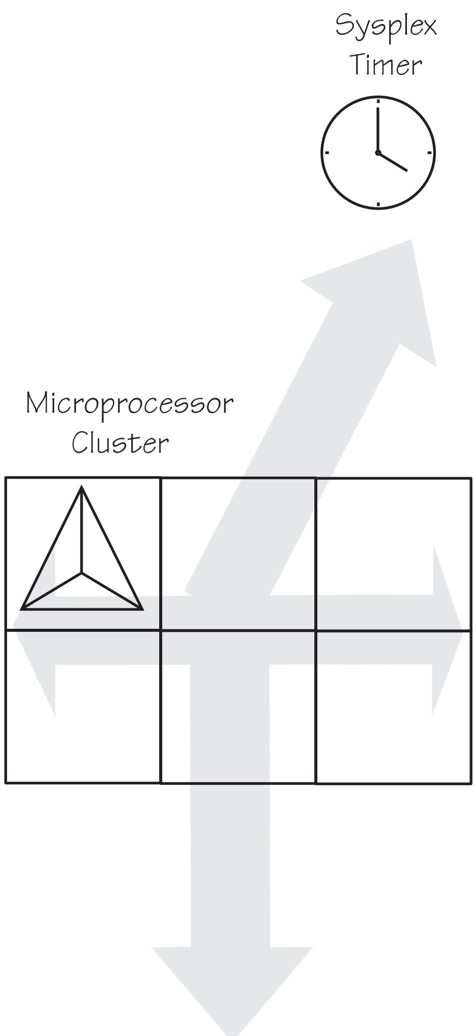

Figure 20. Sysplex with an S/390 9672 Parallel Transaction Server or 9672 Parallel Enterprise Server

The parts of this configuration are:

- S/390 9672 Parallel Transaction Server or 9672 Parallel Enterprise Server with an integrated coupling facility defined in a logical partition
- A Sysplex Timer

Mixed Processor Sysplex

This configuration combines a large processor and an S/390 9672 Parallel Transaction Server or 9672 Parallel Enterprise Server in a sysplex production environment. The processors share S/390 9674 Coupling Facilities C01/C02/C03 Models.

Benefits: In addition to the standard sysplex benefits, such as improved price/performance, greater parallel capacity for a single workload, and environmental savings (space, power, cooling, and weight), this configuration:

- Offers the reliability of a production environment
- Extends sysplex data sharing across different types of processors

- Allows non-disruptive incremental growth, a single CPC at a time

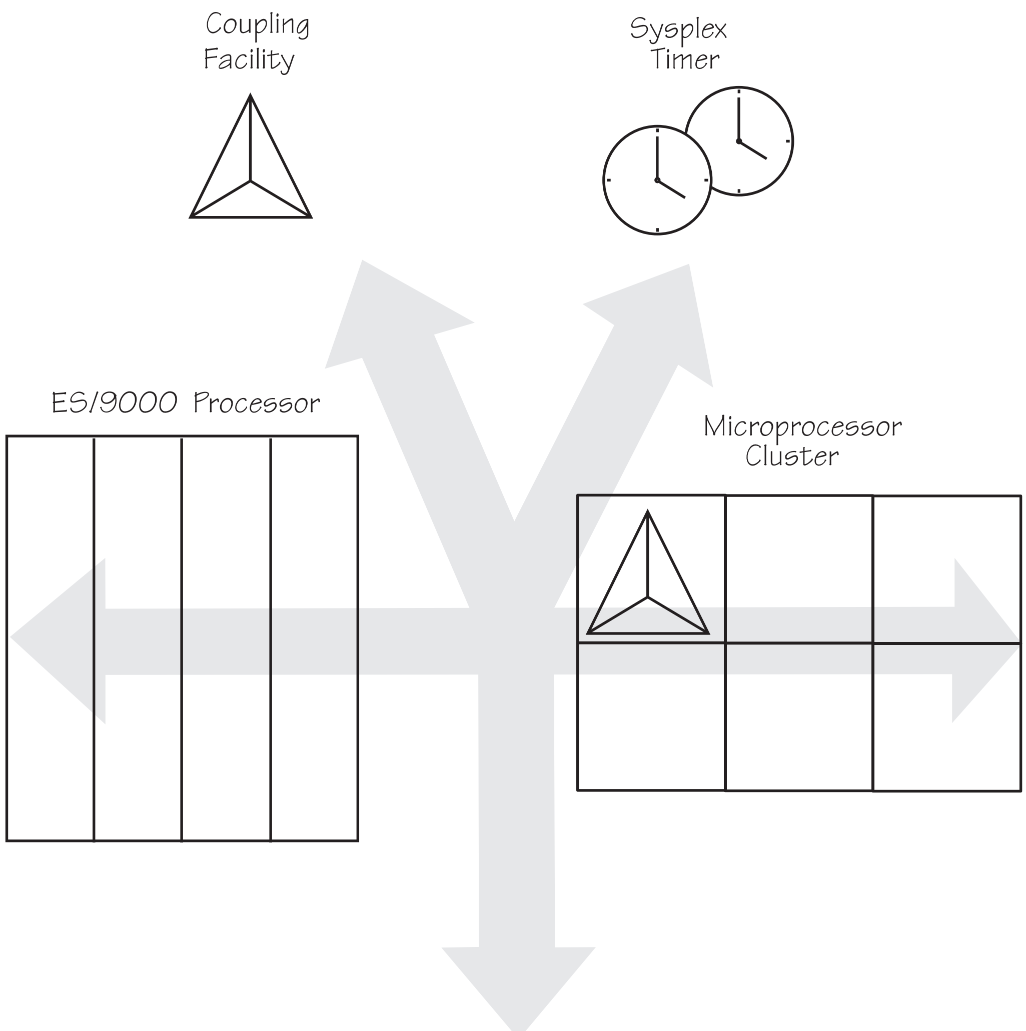

Figure 21. Mixed Processor Sysplex Configuration

The parts of this configuration are:

- ES/9000 large processor, such as:
  – 9021 711-based CPC
  – 9121 511-based CPC
- S/390 9672 Parallel Transaction Server or 9672 Parallel Enterprise Server
- A Sysplex Timer with the high availability feature
- A standalone coupling facility (such as the 9674 Coupling Facility)

Appendix. Notices

This information was developed for products and services offered in the USA.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States and/or other countries:

- Advanced Peer-to-Peer Networking
- APPN
- BookManager
- CBPDO
- CICS
- CICS/ESA
- CICSPlex
- DATABASE 2
- DB2
- DFSMS
- DFSMS/MVS
- Enterprise Systems Connection Architecture
- ES/9000
- ESCON
- IBM
- IBMLink
- IMS
- IMS/ESA
- MVS/ESA
- NetView
- OPC
- OS/2
- OS/390
- Parallel Sysplex
- Processor Resource/Systems Manager
- PR/SM
- RACF
- Resource Measurement Facility
- RMF
- Resource Link
- S/390

- SP
- Sysplex Timer
- System/390
- SystemView
- VisualLift
- VTAM
- z/OS

Other company, product, and service names might be trademarks or service marks of others.

Glossary

Sources of Terms and Definitions

This glossary includes terms and definitions from:

- The *IBM Dictionary of Computing* New York: McGraw-Hill, 1994.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.

Explanation of Cross-References

The following cross-references are used in this glossary:

Contrast with. This refers to a term that has an opposed or substantively different meaning.
See. This refers the reader to multiple-word terms in which this term appears.
See also. This refers the reader to terms that have a related, but not synonymous, meaning.

A

ACDS. Active control data set.

ANO/MVS. Automated Network Operations.

AOC/MVS. Automated Operations Control. The licensed program System Automation for OS/390 includes all of the function previosuly provided by AOC/MVS.

AOR. Application-owning region

APPN. Advanced Peer-to-Peer Networking.

B

basic mode. A central processor mode that does not use logical partitioning. Contrast with *logically partitioned (LPAR) mode*.

batch message processing (BMP) program. An IMS batch processing program that has access to online databases and message queues. BMPs run online, but like programs in a batch environment, they are started with job control language (JCL).

batch-oriented BMP program. A BMP program that has access to online databases and message queues while performing batch-type processing. A batch-oriented BMP does not access the IMS message queues for input or output. It can access online databases, GSAM databases, and MVS files for both input and output.

BMP. Batch message processing (BMP) program.

C

cache structure. A coupling facility structure that enables high-performance sharing of cached data by multisystem applications in a sysplex. Applications can use a cache structure to implement several different types of caching systems, including a store-through or a store-in cache.

cache structure services. MVS services that enable applications in a sysplex to perform operations such as the following on a coupling facility cache structure:

- Manage cache structure resources
- Store data into and retrieve data from a cache structure
- Manage accesses to shared data
- Determine when shared data has been changed
- Determine whether a local copy of shared data is valid.

CBPDO. Custom Built Product Delivery Offering.

CEC. Synonym for central processor complex (CPC).

central processor (CP). The part of the computer that contains the sequencing and processing facilities for instruction execution, initial program load, and other machine operations.

central processor complex (CPC). A physical collection of hardware that includes main storage, one or more central processors, timers, and channels.

CFRM. Coupling facility resource management.

channel-to-channel (CTC). Refers to the communication (transfer of data) between programs on opposite sides of a channel-to-channel adapter (CTCA).

channel-to-channel adapter (CTCA). An input/output device that is used by a program in one system to communicate with a program in another system.

CICS. Customer Information Control System.

CICSplex. A group of connected CICS regions.

CICSPlex SM. CICSPlex System Manager

CMOS. Complementary metal-oxide semiconductor.

COMMDS. Communications data set.

complementary metal-oxide semiconductor (CMOS). A technology that combines the electrical properties of positive and negative voltage requirements to use considerably less power than other types of semiconductors.

couple data set. A data set that is created through the XCF couple data set format utility and, depending on its designated type, is shared by some or all of the MVS systems in a sysplex. See also *sysplex couple data set*.

coupling facility. A special logical partition that provides high-speed caching, list processing, and locking functions in a sysplex.

coupling facility channel. A high bandwidth fiber optic channel that provides the high-speed connectivity required for data sharing between a coupling facility and the central processor complexes directly attached to it.

coupling services. In a sysplex, the functions of XCF that transfer data and status between members of a group residing on one or more MVS systems in the sysplex.

CP. Central processor.

CPC. Central processor complex.

cross-system coupling facility (XCF). XCF is a component of MVS that provides functions to support cooperation between authorized programs running within a sysplex.

CTC. Channel-to-channel.

D

DAE. Dump analysis and elimination.

DASD. Direct access storage device.

data sharing. The ability of concurrent subsystems (such as DB2 or IMS DB) or application programs to directly access and change the same data while maintaining data integrity.

DBCTL. IMS Database Control.

DBRC. Database Recovery Control.

DB2. DATABASE 2 for MVS/ESA.

DB2 data sharing group. A collection of one or more concurrent DB2 subsystems that directly access and change the same data while maintaining data integrity.

DB2 PM. DB2 Performance Monitor.

DFSMS. Data Facility Storage Management Subsystem.

dpAM. IBM SystemView Data Processing Accounting Manager/MVS.

E

EMIF. ESCON Multiple Image Facility.

Enterprise Systems Connection (ESCON). A set of products and services that provides a dynamically connected environment using optical cables as a transmission medium.

EPDM. IBM SystemView Enterprise Performance Data Manager/MVS.

ESCD. ESCON Director.

ESCM. ESCON Manager. The licensed program System Automation for OS/390 includes all of the function previosuly provided by ESCM.

ESCON. Enterprise Systems Connection.

ETR. External Time Reference. See also *Sysplex Timer*.

F

FOR. File-owning region.

frame. For a System/390 microprocessor cluster, a frame contains one or two central processor complexes (CPCs), support elements, and AC power distribution.

G

global resource serialization. A function that provides an MVS serialization mechanism for resources (typically data sets) across multiple MVS images.

global resource serialization complex. One or more MVS systems that use global resource serialization to serialize access to shared resources (such as data sets on shared DASD volumes).

GSAM. Generalized Sequential Access Method.

GTF. Generalized trace facility.

H

Hardware Management Console. A console used to monitor and control hardware such as the System/390 microprocessors.

HCD. Hardware Configuration Definition.

highly parallel. Refers to multiple systems operating in parallel, each of which can have multiple processors. See also *n-way*.

I

ICMF. Integrated Coupling Migration Facility.

image server. A high-capacity optical storage device or a computer that each computer and image workstation on a network can use to access and retrieve image objects that can be shared among the attached computers and image workstations.

IMS. Information Management System.

IMS DB. Information Management System Database Manager.

IMS DB data sharing group. A collection of one or more concurrent IMS DB subsystems that directly access and change the same data while maintaining data integrity.

IMS TM. Information Management System Transaction Manager.

in-doubt period. The period during which a unit of work is pending during commit processing that involves two or more subsystems. See also *in-doubt work unit*.

in-doubt work unit. In CICS/ESA and IMS/ESA, a piece of work that is pending during commit processing; if commit processing fails between the polling of subsystems and the decision to execute the commit, recovery processing must resolve the status of any work unit that is in doubt.

integrated operations workstation. A programmable workstation (PWS) from which an individual can access multiple products to perform a set of tasks, in some cases without knowing which particular product performs a specific task.

IOCDS. Input/output configuration data set.

IOCP. Input/output configuration program.

IODF. Input/output definition file.

IRLM. Internal resource lock manager.

ISPF. Interactive System Productivity Facility.

J

JES2. Job Entry Subsystem 2.

JES3. Job Entry Subsystem 3.

L

LIC. Licensed Internal Code.

list structure. A coupling facility structure that enables multisystem applications in a sysplex to share information organized as a set of lists or queues. A list structure consists of a set of lists and an optional lock table, which can be used for serializing resources in the list structure. Each list consists of a queue of list entries.

list structure services. MVS services that enable multisystem applications in a sysplex to perform operations such as the following on a coupling facility list structure:
- Read, update, create, delete, and move list entries in a list structure
- Perform serialized updates on multiple list entries in a list structure
- Monitor lists in a list structure for transitions from empty to non-empty.

lock structure. A coupling facility structure that enables applications in a sysplex to implement customized locking protocols for serialization of application-defined resources. The lock structure supports shared, exclusive, and application-defined lock states, as well as generalized contention management and recovery protocols.

lock structure services. MVS services that enable applications in a sysplex to perform operations such as the following on a coupling facility lock structure:
- Request ownership of a lock
- Change the type of ownership for a lock
- Release ownership of a lock
- Manage contention for a lock
- Recover a lock held by a failed application.

logical partition (LP). A subset of the processor hardware that is defined to support an operating system. See also *logically partitioned (LPAR) mode*.

logically partitioned (LPAR) mode. A central processor complex (CPC) power-on reset mode that enables use of the PR/SM feature and allows an operator to allocate CPC hardware resources (including central processors, central storage, expanded storage, and channel paths) among logical partitions. Contrast with *basic mode*.

loosely coupled. A multisystem structure that requires a low degree of interaction and cooperation between multiple MVS images to process a workload. See also *tightly coupled*.

LP. Logical partition.

LPAR. Logically partitioned (mode).

M

m-image. The number (m) of MVS images in a sysplex. See also *n-way*.

massively parallel. Refers to thousands of processors in a parallel arrangement.

member. A specific function (one or more modules/routines) of a multisystem application that is defined to XCF and assigned to a group by the multisystem application. A member resides on one system in the sysplex and can use XCF services to communicate (send and receive data) with other members of the same group.

microprocessor. A processor implemented on one or a small number of chips.

mixed complex. A global resource serialization complex in which one or more of the systems in the global resource serialization complex are not part of a multisystem sysplex.

MP. Multiprocessor.

MRO. Multiregion operation.

MSC. Multiple Systems Coupling.

multi-MVS environment. An environment that supports more than one MVS image. See also *MVS image* and *sysplex*.

Multiple Systems Coupling (MSC). An IMS facility that permits geographically dispersed IMS subsystems to communicate with each other.

multiprocessing. The simultaneous execution of two or more computer programs or sequences of instructions. See also *parallel processing*.

multiprocessor (MP). A CPC that can be physically partitioned to form two operating processor complexes.

multisystem application. An application program that has various functions distributed across MVS images in a multisystem environment.

multisystem environment. An environment in which two or more MVS images reside in one or more processors, and programs on one image can communicate with programs on the other images.

multisystem sysplex. A sysplex in which two or more MVS images are allowed to be initialized as part of the sysplex. See also *single-system sysplex*.

MVS image. A single occurrence of the MVS/ESA operating system that has the ability to process work.

MVS system. An MVS image together with its associated hardware, which collectively are often referred to simply as a system, or MVS system.

MVS/ESA. Multiple Virtual Storage/ESA.

MVSCP. MVS configuration program.

N

n-way. The number (*n*) of CPs in a CPC. For example, a 6-way CPC contains six CPs.

NJE. Network job entry.

O

OLTP. Online transaction processing.

OPC/ESA. Operations Planning and Control.

operating system (OS). Software that controls the execution of programs and that may provide services such as resource allocation, scheduling, input/output control, and data management. Although operating systems are predominantly software, partial hardware implementations are possible. (T)

P

parallel processing. The simultaneous processing of units of work by many servers. The units of work can be either transactions or subdivisions of large units of work (batch). See also *highly parallel*.

Parallel Sysplex. A sysplex that uses one or more coupling facilities.

partitionable CPC. A CPC that can be divided into 2 independent CPCs. See also *physical partition, single-image mode, MP, side*.

physical partition. Part of a CPC that operates as a CPC in its own right, with its own copy of the operating system.

physically partitioned (PP) configuration. A system configuration that allows the processor controller to use both central processor complex (CPC) sides as individual CPCs. The A-side of the processor controller controls side 0; the B-side of the processor controller controls side 1. Contrast with *single-image (SI) configuration*.

PR/SM. Processor Resource/Systems Manager.

processor controller. Hardware that provides support and diagnostic functions for the central processors.

Processor Resource/Systems Manager (PR/SM). The feature that allows the processor to use several MVS images simultaneously and provides logical partitioning capability. See also *LPAR*.

Q

QOR. Queue-owning region.

R

RACF. Resource Access Control Facility.

RMF. Resource Measurement Facility.

S

SCDS. Source control data set.

SDSF. System Display and Search Facility.

SEC. System Engineering Change.

serialized list structure. A coupling facility list structure with a lock table containing an array of exclusive locks whose purpose and scope are application-defined. Applications can use the lock table to serialize on parts of the list structure, or resources outside the list structure.

side. A part of a partitionable CPC that can run as a physical partition and is typically referred to as the A-side or the B-side.

single point of control. The characteristic a sysplex displays when you can accomplish a given set of tasks from a single workstation, even if you need multiple IBM and vendor products to accomplish that particular set of tasks.

single system image. The characteristic a product displays when multiple images of the product can be viewed and managed as one image.

single-image (SI) mode. A mode of operation for a multiprocessor (MP) system that allows it to function as one CPC. By definition, a uniprocessor (UP) operates in single-image mode. Contrast with *physically partitioned (PP) configuration*.

single-MVS environment. An environment that supports one MVS image. See also *MVS image*.

single-system sysplex. A sysplex in which only one MVS system is allowed to be initialized as part of the sysplex. In a single-system sysplex, XCF provides XCF services on the system but does not provide signalling services between MVS systems. See also *multisystem sysplex, XCF-local mode*.

SLR. Service Level Reporter.

SMF. System management facilities.

SMP/E. System Modification Program Extended.

SMS. Storage Management Subsystem.

SMS communication data set. The primary means of communication among systems governed by a single SMS configuration. The SMS communication data set (COMMDS) is a VSAM linear data set that contains the current utilization statistics for each system-managed volume, which SMS uses to help balance space usage among systems.

SMS configuration. The SMS definitions and routines that the Storage Management Subsystem uses to manage storage.

SMS system group. All systems in a sysplex that share the same SMS configuration and communications data sets, minus any systems in the sysplex that are defined individually in the SMS configuration.

structure. A construct used by MVS to map and manage storage on a coupling facility. See *cache structure, list structure*, and *lock structure*.

support element. A hardware unit that provides communications, monitoring, and diagnostic functions to a central processor complex (CPC).

symmetry. The characteristic of a sysplex where all systems, or certain subsets of the systems, have the same hardware and software configurations and share the same resources.

SYSLOG. System log

sysplex. A set of MVS systems communicating and cooperating with each other through certain multisystem hardware components and software services to process customer workloads. See also *MVS system, Parallel Sysplex*.

sysplex couple data set. A couple data set that contains sysplex-wide data about systems, groups, and members that use XCF services. All MVS systems in a sysplex must have connectivity to the sysplex couple data set. See also *couple data set*.

Sysplex Timer. An IBM unit that synchronizes the time-of-day (TOD) clocks in multiple processors or processor sides. External Time Reference (ETR) is the MVS generic name for the IBM Sysplex Timer (9037).

system control element (SCE). Hardware that handles the transfer of data and control information associated with storage requests between the elements of the processor.

System/390 microprocessor cluster. A configuration that consists of central processor complexes (CPCs) and may have one or more coupling facilities.

T

tightly coupled. Multiple CPs that share storage and are controlled by a single copy of MVS. See also *loosely coupled, tightly coupled multiprocessor*.

tightly coupled multiprocessor. Any CPC with multiple CPs.

TOR. Terminal-owning region.

TSCF. Target System Control Facility. The licensed program System Automation for OS/390 includes all of the function previosuly provided by TSCF.

U

uniprocessor (UP). A CPC that contains one CP and is not partitionable.

UP. Uniprocessor.

V

VM. Virtual Machine.

VSAM. Virtual Storage Access Method.

VTAM. Virtual Telecommunications Access Method.

W

WLM. MVS workload management.

X

XCF. Cross-system coupling facility.

XCF PR/SM policy. In a multisystem sysplex on PR/SM, the actions that XCF takes when one MVS system in the sysplex fails. This policy provides high availability for multisystem applications in the sysplex.

XCF-local mode. The state of a system in which XCF provides limited services on one system and does not provide signalling services between MVS systems. See also *single-system sysplex*.

XRF. Extended recovery facility.

Index

Numerics

A

B

C

D

E

F

H

I

Readers' Comments — We'd Like to Hear from You

z/OS
Parallel Sysplex Overview:
Introducing Data Sharing and
Parallelism in a Sysplex

Publication No. SA22-7661-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.

Readers' Comments — We'd Like to Hear from You
SA22-7661-00

Cut or Fold
Along Line

Fold and Tape **Please do not staple** Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY
12601-5400

Fold and Tape **Please do not staple** Fold and Tape

Cut or Fold
Along Line

Program Number: 5694-A01

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SA22-7661-00